AD-A086 490 GEORGIA INST OF TECH ATLANTA ENGINEERING EXPERIMENT —ETC F/G 9/2
NAVY EMBEDDED COMPUTER ACCREDITATION PROGRAM.(U)
APR 80 B B WISE N00014-79-C-0722
UNCLASSIFIED GIT/EES-A-2486 NL

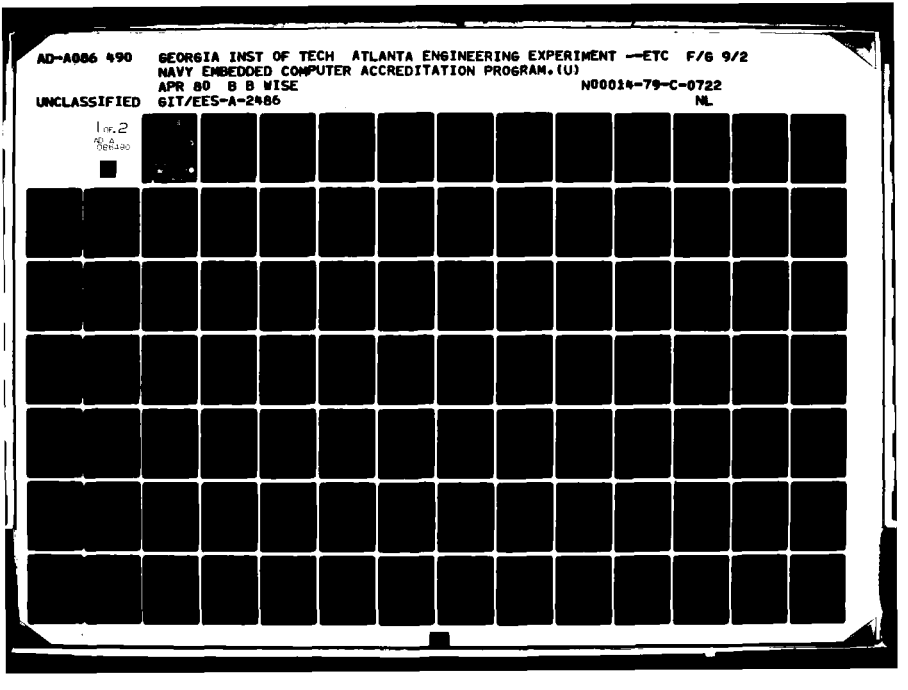

FINAL TECHNICAL REPORT
GIT/EES PROJECT A-2486

LEVEL II 12

NAVY EMBEDDED COMPUTER ACCREDITATION PROGRAM

By
Billy B. Wise
Computer Science and Technology Laboratory

Prepared for
DEPARTMENT OF THE NAVY
Office of Assistant Secretary of the Navy
(Research, Engineering and Systems)

Under
Contract N00014-79-C-0722 new

April 1980

GEORGIA INSTITUTE OF TECHNOLOGY
Engineering Experiment Station
Atlanta, Georgia 30332

GTEES

GEORGIA INSTITUTE OF TECHNOLOGY
Engineering Experiment Station
Atlanta, Georgia 30332

NAVY EMBEDDED COMPUTER ACCREDITATION PROGRAM

Final Technical Report

by

Billy B. Wise

of the

Computer Science and Technology Laboratory

GIT/EES Project A-2486

H. Bennett Teates - Project Director

April 1980

Prepared for

DEPARTMENT OF THE NAVY
Office of Assistant Secretary of the Navy
(Research, Engineering and Systems)

through

The Office of Naval Research
under
Contract N00014-79-C-0722

ACKNOWLEDGEMENTS

This report represents the cumulative effort of the several individuals who contributed to the project. In particular, it is worthy to note the contributions of John F. Passafiume, who contributed his knowledge about the Army's Military Computer Family (MCF) Program; Douglas E. Wrege, who contributed materially to the questionnaire and basic technical issues; Fred L. Cox, who contributed to the questionnaire and added his in-depth knowledge of Ada and its future; and to Harold S. Stone, who, under subcontract to the Engineering Experiment Station, made significant contributions to the life cycle cost issues. A special note of recognition is deserved by Billy B. Wise, who authored the final report and presented project results to the Navy's review committee.

H. Bennett Teates

Project Director

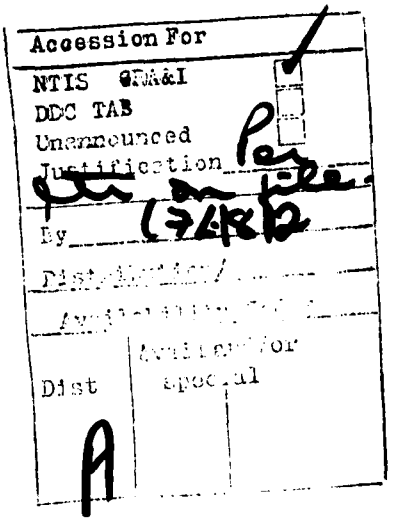

ABSTRACT

The Navy Embedded Computer Accreditation Program is the examination of the viability and strategy appropriate to a new, more nearly "optimal" acquisition policy for embedded computers. Under accreditation, the Navy would approve (accredit) a controlled number of computers for use by project managers as the computers become available and meet certain qualification criteria relating to life cycle cost, performance and logistic supportability. This report addresses the issues and establishes a tentative set of accreditation criteria which provide a means for the Navy to move smoothly from the current policy of standardization on a single computer in a performance range to the accreditation policy.

Key Words

Standardization

Accreditation

Embedded Computers

Logistics Cost

Life Cycle cost

EXECUTIVE SUMMARY

The Navy's policy on the acquisition of computers for embedded applications has been to standardize on one computer in a performance range, primarily for purposes of hardware maintainability at sea. The price of this standardization has been loss of continuing competition, lack of timely technology infusion, and costly computers far behind the state of the art. Recent emphasis on these detrimental aspects has prompted Navy interest in moving away from the present acquisition policy and toward a less rigid acquisition strategy called accreditation. In defining an accreditation program, there are two problems which must be considered:

(1) How to move away from hardware standardization on a single computer in each performance range.
(2) How to reduce the required numbers of computer maintenance technicians on Navy ships at sea.

The purpose of the study reported in this document was to examine the viability and strategy appropriate to the new, more nearly "optimal" acquisition policy of accreditation. Under accreditation the Navy would approve (accredit) a controlled number of computers for use by project managers as the computers become available and meet certain qualification criteria relating to life-cycle cost, performance and logistic supportability. The objectives of the study were to:

(1) Design an accreditation program (i.e., determine essential criteria appropriate to accreditation program implementation).
(2) Investigate maintenance schemes to reduce required numbers of maintenance technicians.
(3) Prepare a transition strategy to move from standardization to accreditation.

In designing an accreditation program, it was necessary to establish some criteria against which candidate computers could be evaluated. The Army's Military Computer Family, the Navy Embedded Computer System, and the Air Force Electronics Standardization Program, among others, were examined to assemble a list of potential accreditation criteria for consideration. Life-cycle cost factors and the rate of technological advances in the computer industry were examined to determine their effect upon accreditation cycle length, defined as the period of time between opportunities for computer manufacturers to offer their machines for accreditation. Consideration was also given to maintenance techniques appropriate for reducing the requirement for human participation in computer maintenance at sea. A questionnaire regarding the proposed accreditation program was prepared and sent to selected computer manufacturing firms to sample industry response to the proposal and to solicit suggestions.

A performance matrix, defined by five application classes and three performance levels, was devised to allow classification of accredited computers. Each cell of the matrix defines an accreditation list and each list may have multiple entries (Figure 1). The collection of recommended accreditation criteria is presented in chart format (Figure 2). On the left side are displayed interim criteria, those intended for application at program initiation. In the middle are the mid-term criteria, to be applied at the beginning of the second accreditation cycle. On the right side are listed criteria associated with the target accreditation program, the recommended mature form for the program. The accreditation criteria are divided into three groups depending upon their functions, which are to define mandatory features, to classify as to performance, and to rank according to life-cycle cost (LCC). The criteria associated with mandatory features are pass/fail criteria. The performance classification criterion will place an offered

Figure 1

PERFORMANCE BY APPLICATION CLASS MATRIX

PERFORMANCE LEVEL (KOPs)	Application Class				
	C^2	Comm. Switch	Number Processor	DBM	Character Processor
HIGH					
MEDIUM					
LOW					

Interim	Mid Term	Target
	Mandatory Requirements	
Emulate ISAs of current standard computers	Emulate ISAs of current standard computers	ISA standardization at a common HOL level
Use current standard language	Use Ada HOL	Use Ada HOL
Validation of hardware compliance to emulated ISA standard	Validation of hardware compliance to emulated ISA standard	Validation of hardware compliance to Ada HOL/ISA standard
Require use of SEMs	Standardize at box level for accreditation	Standardize at box level for accreditation
Standard Interface between boxes	Standard Interface between boxes	Standard Interface between boxes
Two Thousand hour MTBF	Three Thousand hour MTBF	Four Thousand hour MTBF
Require use of SEMs	Require BIT to diagnose to module level	Require BIT to diagnose to module level
Require use of SEMs	Maintain and Spare at Module Level	Maintain and Spare at Module Level
	Performance Classification	
Use existing performance levels	Use accreditation performance matrix	Use accreditation performance matrix
	Ranking	
Consider LCC in evaluating candidates	LCC Model is major discriminator between candidates	LCC Model is major discriminator between candidates
	Accreditation Cycle Length	
Five years	Five years	Five Years
	Number of Computers per List	
Two in each performance range	Competition may limit number on list	Competition may limit number on list

machine in its appropriate cell in the performance matrix, i.e., on its appropriate accreditation list. The LCC ranking criterion will determine the offered machine's rank within the accreditation list.

A milestone plan was drawn to show the time schedule and steps to proceed from standardization to accreditation (Figure 3). The Government will have to take a strong leadership position in defining and supporting the concept of accreditation if it is to become generally accepted by both military project managers and the computer industry.

Figure 3
TRANSITION PLAN
MILESTONES

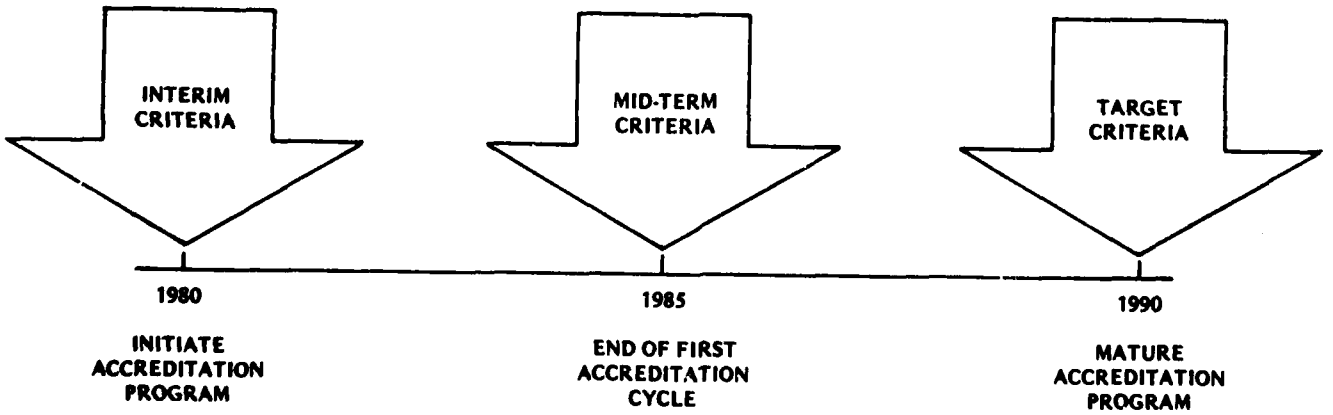

The iterim criteria listed in Figure 2 can be implemented immediately. However, implementation of the mid-term and target criteria will require two types of additional effort. One type is the additional study of accreditation criteria and specifications. Included in this group are the following:

- o A detailed definition of application classes for Navy systems.
- o Development of benchmark routines or instruction mix equations for application classes.
- o Development of a comprehensive life-cycle cost model.
- o Determination of the upper limit on accredited machines.

The other type of effort required deals with Navy policy level emphasis on DoD programs and the implementation of Navy-wide efforts to accumulate and codify requisite data and to initiate programs to take advantage of the accreditation strategy. Included in this group are the following:

- o Support of DoD efforts in the Ada, VLSI and VHSIC programs.
- o Provide firm guidance to industry regarding requirement to develop and implement built-in-test and fault tolerant machines.
- o Development of a plan and initiation of a program to collect data in support of the life-cycle cost model elements.
- o Modify maintenance and training policy in accordance with the accreditation program.

Finally, the Navy must take an active leadership role. It must establish accreditation as its policy, educate its project managers as to its benefits, announce its plans and goals to the industrial community, and move to support the program through its fruition.

TABLE OF CONTENTS

NAVY EMBEDDED COMPUTER ACCREDITATION PROGRAM

(NECAP)

1.0 INTRODUCTION

1.1 Background

Embedded computers perform vital functions in assisting Navy combat units to conduct their warfare missions in the face of an increasing hostile threat capability. Without these computers, the offensive and defensive weapons systems, the navigation, and command and control systems in Navy platforms could not perform at the required level to assure survivability of the combat unit and maximize its war-fighting effectiveness. It has been estimated that by 1985 the number of embedded computers in use throughout the Navy will be over three times the number currently in use. One of the unique fundamentals of Navy use of and plans for embedded computers is the requirement for these machines to operate at sea. This requirement casts a heavy burden on the inherent reliability needed in a computer, the spare-parts philosophy, the technology of maintenance, and on the availability and quality of Navy maintenance personnel at sea.

Possible acquisition policies for Navy embedded computers range from absolute standardization on a single computer for all applications to completely unconstrained development and acquisition of computers by each Navy project manager according to his own requirements. Operating at either extreme of this range is currently undesirable and some optimal acquisition strategy lies somewhere in between these two end points. The Navy acquisition policy currently in effect consists of standardization on a single computer in a given performance range for use in embedded applications aboard ship.

The predominate reason for the present Navy policy of standardization is hardware maintainability.[1] Concerns for hardware maintenance and spares supplies on ships at sea impose a need for some level of standardization on hardware. This existing policy of hardware standardization has contributed toward alleviating the problems of sparing, maintenance training, and overall operability at sea.

1.2 Definition of the Problem

1.2.1 Detrimental Aspects of Hardware Standardization. Hardware standardization yields benefits in the areas of sparing, maintenance, and operability; however, there is a penalty associated with the attainment of these benefits. The price of standardizing on one computer in a performance range is loss of continuing competition, lack of timely technology infusion into embedded computer systems, and costly computers far behind commercial state of the art. Recent emphasis on these detrimental aspects of hardware standardization has prompted Navy interest in moving away from the present acquisition policy.

1.2.2 Availability of Maintenance Personnel. A second problem is the availability of maintenance personnel for embedded computers at sea. Increasing use of high-technology and sophisticated systems by the Navy has generated an ever-growing need for greater numbers of highly trained maintenance technicians. The Navy has taken measures to provide additional personnel to meet the near-term maintenance requirements caused by the influx of embedded tactical computers. However, in the longer term there are indications that the projected high acquisition levels of embedded tactical

computers, coupled with a perceived reduction in size of the available national manpower pool from which maintenance trainees are recruited, are likely to result in a growing shortfall in numbers of trained embedded computer maintenance technicians.

The two defined problems which the Navy faces in embedded computers are:

(1) How to move away from hardware standardization on a single computer in each performance range.

(2) How to reduce the required numbers of computer maintenance technicians on Navy ships at sea.

It is readily apparent that these two problems are interrelated. As stated previously, the reason for the present Navy policy of standardization is hardware maintainability. Thus, any move away from this policy of standardization toward a more flexible acquisition policy allowing procurement of more types of embedded computers will adversely affect the maintenance problem, tending to increase the numbers of maintenance technicians required and the level of maintenance training required. This increased maintenance requirement will further aggravate the problem of projected maintenance technician shortfalls.

1.3 Purpose of the Study

The purpose of this study is to determine the characteristics of an "optimal" acquisition strategy, as referred to in paragraph 1.1. This so-called optimal strategy will be less rigid than the current policy of strict standardization on one specific computer in each performance range. Under this less rigid policy the Navy would approve (accredit) a controlled number of computers for use by project managers, as the hardware becomes available from industry and meets certain qualification criteria relating to life-cycle cost,

performance, and logistic supportability. This less rigid, optimal acquisition strategy is referred to as accreditation.

The basic goal of accreditation is to permit the Navy to obtain the benefits of new computer technology and open competition as frequently as possible, while at the same time satisfying operational requirements, military logistics, and budget constraints. The accreditation approach to embedded computer acquisition shows potential to accomplish the following desirable goals:

a. Stimulate competition - the present standardization policy virtually amounts to sole source procurement. Industry is not stimulated to provide better equipment at lower cost.

b. Ease technology insertion - the present standardization policy thwarts the injection of new technology, which has advanced considerably in the years since the policy was established. For example, it is now possible to acquire a computer containing approximately 20 cards with comparable performance to that of a 1-bay UYK-7 computer containing 800 cards.

c. Increase flexibility of choice to project managers - the present standardization policy gives project managers no choice of selection. Accreditation would offer at least some limited choices and is closer to the total-system concept in project development.

d. Shorten acquisition cycle - the acquisition cycle for embedded computers ranges from five to eight years, with seven years as an average. Accreditation has the potential to significantly shorten the time required for acquisition.

There are three specific goals of this study effort:

1. Design an accreditation program

2. Investigate maintenance schemes to reduce required numbers of maintenance technicians

3. Prepare a transition plan to move from standardization to accreditation

1.3.1 Design an Accreditation Program. The accreditation program design will include a target accreditation goal, i.e., the final form of an accreditation program to be attained after a suitable period of transition; a set of recommended accreditation criteria by which candidate computers will be evaluated; and a procedure for conducting the accrediting process on some established schedule.

1.3.2 Investigate Maintenance Schemes to Reduce Numbers of Maintenance Technicians Required. The intent of this goal will be the examination of maintenance philosophies and techniques which will reduce the need for human participation in the troubleshooting and repair of embedded computers. Certain maintenance techniques, such as built-in-test, will surely impinge upon accreditation criteria.

1.3.3 Prepare a Transition Plan. The existing investment in application software and software tools, among other factors, will not permit an instantaneous shift away from the current acquisition policy of standardization to the proposed policy of accreditation. Certain necessary accreditation criteria, to assure operability of an accreditation acquisition scheme, cannot be attained without advances in computer technology. For these reasons, a transition plan will be formulated to provide guidance for an

orderly transition from standardization to accreditation. This plan will include projected milestones for attainment of necessary technological advances.

2.0 METHODOLOGY

2.1 Selection of Candidate Accreditation Criteria

There have been numerous study efforts sponsored by the various Services in recent years regarding the problems of hardware and software proliferation, acquisition, and support. The approach here was to examine a number of these programs and, drawing upon the successes and failures of these efforts, assemble a list of candidate accreditation criteria. Among the programs examined were the Army's Military Computer Family, the Navy Embedded Computer System, and the Air Force Electronics Standardization Program.

2.2 Accreditation Cycle Length

One of the goals in moving away from standardization toward accreditation as an acquisition policy is to reduce acquisition time. The implication is that the accreditation cycle length should be shorter than the present average seven years required to acquire a computer under standardization. Two factors which may have an effect on the determination of accreditation cycle length are the span of time between significant advances in computer hardware technology and between major developments in Instruction Set Architectures (ISA). There are also considerations in the determination of accreditation cycle length from a life cycle cost standpoint. On the logistic side, the preference would seem to be toward longer cycle length, because of the investment in an inventory system, spare parts, and training. The approach here was to find an area of mutual overlap between technological generations and life cycle savings accruing to injection of the new technology.

2.3 Maintenance Schemes

It has been pointed out that projected shortages of available computer maintenance personnel will only be aggravated by any move away from the current policy of hardware standardization. The approach here was to examine specific ways of reducing the requirement for human participation in embedded computer maintenance.

2.4 Comments and Projections from Industry

After assessing the trends in technology, using the lessons learned from related programs to establish some tentative accreditation criteria, and determining deficiencies in life cycle cost data; a questionnaire was prepared and sent to selected industrial firms. (See Appendix A for a copy of the questionnaire and list of addresses). This approach provided a means by which we could substantiate our preliminary findings and judgments, obtain new ideas and approaches, and determine the general acceptance of the accreditation concept.

Figure 1

METHODOLOGY

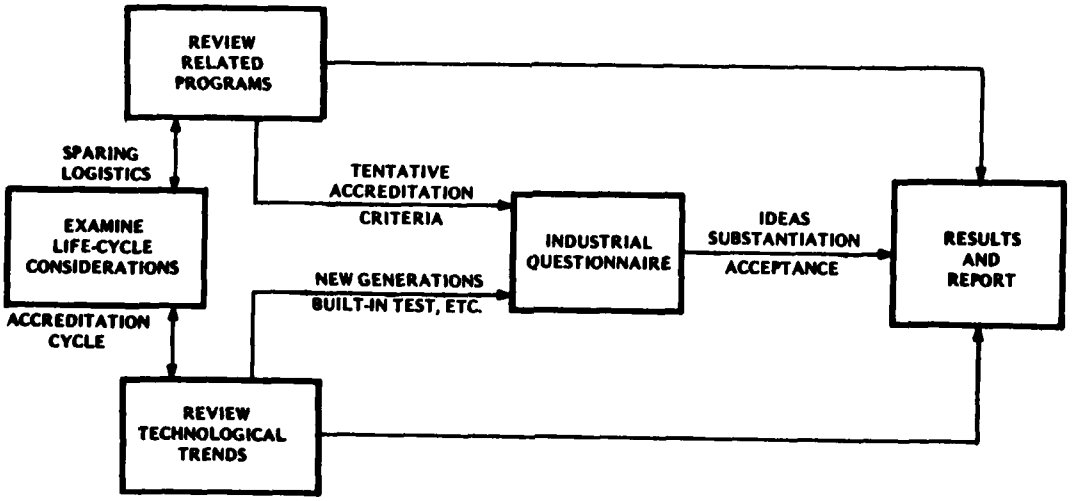

3.0 THE ACCREDITATION PROGRAM

The idea behind accreditation is that computer vendors would be invited on a periodic basis to compete their products against a set of established accreditation criteria. A certain number of computing machines which meet these established criteria would be placed on an approved list for use by program managers in selecting computers for use in Navy systems. The result would be a readily available pool of computing machines, having the latest technological improvements, from which to select. This scheme has the potential to reduce significantly the acquisition time for embedded computers, to assure availability of the latest technology, and to allow for more competition in the marketplace.

There are several factors which will require close attention if such an acquisition policy is to be implemented successfully. Appropriate criteria for accreditation will have to be developed to ensure that user requirements can be fulfilled by the accredited machines and to allow for meaningful competition between computer vendors. A determination will have to be made regarding the appropriate number of machines to be placed on the accredited list. The problems associated with the current standardization policy imply that one computer for each performance range is marginally adequate. Conversely, because of the somewhat unique Navy problems associated with hardware maintainability at sea, there is a need to keep the number of accredited computers on the list from growing too large. The length of the accreditation cycle must be such that it encourages competition among vendors, while also taking into consideration effects upon logistics costs to the Navy. Finally, consideration must be given to criteria for removal from an accredited list. Technological advances may be the most important factor

here; however, in order to keep competition alive, some other criteria for removal must be established in the event that technological development or user needs for new technology diminish.

3.1 Accreditation Criteria

This section contains discussion of candidate accreditation criteria, including the comments obtained from industrial firms via the questionnaire. (See Appendix B for a summary of survey results.) Among the topics discussed are Instruction Set Architecture (ISA) and High Order Language (HOL) standardization, performance factors, performance levels, standard interface, level of hardware standardization, and life-cycle cost (LCC).

3.1.1 ISA/HOL Standardization

3.1.1.1 Background. An Instruction Set Architecture (ISA) is defined to be all of the timing independent information about a computer necessary to write software for that machine. An ISA standard does not include instruction timing information nor any implementation details not visible to the programmer (e.g., the existence of cache memory, number of memory parity bits, add time, multiply time, interrupt latency, etc.). It does include information (e.g., privileged instructions and memory translation instructions) necessary to implement operating systems and system software.

It is useful to decompose the structure of a computer system into a series of levels, ranging from the hardware circuit level to the ISA that the computer programmer sees. For the purpose of development of this concept, it will be assumed that the computer is microprogrammed with the microprocessor engine labeled as the level one machine. This level-one ISA is used to

implement a level-two ISA through a microcode interpreter. This level-two ISA is what is commonly referred to as the conventional machine language ISA.

In general, new conceptual machines or ISA levels can be created from the previous level by one of two processes, interpretation or translation. The process of translation refers to replacing each instruction of ISA(n) with an equivalent sequence of instructions in ISA(n-1). The result is a program consisting entirely of ISA(n-1) instructions, which the underlying level machine may then execute. The process of interpretation involves implementing a program in ISA(n-1) which takes instructions in ISA(n) as input data and executes them by examining each instruction in turn and executing the equivalent sequence of ISA(n-1) instructions directly.

3.1.1.2 Issues. Studies in software engineering suggest that ISA standardization should be at the High Order Language (HOL) machine level. Such a standardization policy would result in maximizing the robustness of software systems and allow the freedoms desired for technology infusion. The standardization of HOL and instruction set architecture allows the standardization of the compiler and reuse of associated HOL source code and instruction set code software program modules. This policy in turn provides a relatively easy way to achieve and ensure future software cost reductions. [2] In a study for the Army, Stone reported that adoption of a single ISA standard has the potential to yield nearly a fifty per cent life cycle cost savings compared to a multiple ISA situation.[3] The recommended target accreditation factor is ISA standardization at a common HOL level. This factor would allow suppliers to qualify for an accreditation list through a variety of implementation levels. For example, they could supply a machine with an underlying ISA plus the necessary compilers, run-time systems, etc., to

implement the HOL/ISA, or implement as much of the HOL/ISA at the level-two machine as technology would permit. While standardization at the HOL level is for many reasons impractical at the current time, it is considered to be the desired long-range goal.

In the survey questionnaire, respondents were asked their opinion as to whether HOL ISA standardization is a worthwhile and realizable goal. With only minor caveats, the response was unanimously in the affirmative.

It is recognized that there are certain requirements that must be met before HOL/ISA standardization is reasonable. Among these is the existence of an HOL with attributes which will provide freedom from dependence on traditional machine language programming. That the DoD is moving in this direction is evidenced by DoD Instruction 5000.31 and the current Ada language effort.

DoD Instruction 5000.31 significantly reduced the number of programming languages approved for use in new systems. The DoD Common High Order Language program was initiated in 1975 with the goal of establishing a single high order, machine independent language for new DoD embedded computer systems. This language, called Ada, is optimized for use in and development of embedded computer systems, and by its design is to substantially reduce the need for and use of machine language programming. Ada is machine independent, thereby achieving true transportability of software developed using the language. The major recognized benefits of a common high order language are derived from Ada's appropriateness to military applications, from the portability that comes with a machine independent language, from the availability of software resulting from acceptance of the language for nonmilitary applications, and most importantly, from the use of Ada as a mechanism for introducing and distributing effective software development and support environments to firms

developing and evolving military systems.

Sixty-six per cent of the respondents to the questionnaire agreed that Ada would be an appropriate HOL. Although the response was quite favorable toward Ada, several firms also wanted additional accredited HOLs, such as F77, FIV, CMS-2, J73, J731, and ATLAS. Given the profit-oriented nature of the business community, this desired diversity is understandable. However, such diversity would severely detract from the economy resulting to the DoD as the user of a single standard HOL.

Several questions were included in the computer industry survey regarding methods of progressing toward an HOL/ISA standard. To a question regarding the pace of movement toward the use of HOLs, responses were mixed. There were an equal number of yes and no answers, but the rationale behind the negative responses implicate a lack of direction or leadership. In very general terms, there was a feeling of dissatisfaction with existing government directives relating to the use of HOLs; but there was no consensus as to what should be done -- only that more decisive direction is required.

Industry also was asked what technical problems needed to be solved before an HOL/ISA standardization would be reasonable. Among the problems perceived were the following:

1. metrics
2. testing techniques
3. transportability to host computers
4. application to process of designing application programs
5. access to software tools on a tri-service basis
6. compiler availability
7. compiler must be in the public domain
8. machine architecture complimentary to HOL
9. definition of minimum hardware required to support HOL standards
10. definition of allowable host machines

It would appear that none of these perceived problems is insolvable. One and Two should be addressed by an extension of the Ada Compiler Validation Capability currently being developed under DARPA support. Three, Four, and Five should be solved by the inherent nature and the hoped-for global adoption of the Ada language. The Air Force and the Army currently have programs to develop Ada compilers and software tools which will go into the public domain following the validation process. Eight, Nine, and Ten should be solved in the long term through development by government and/or industry of an Ada machine.

An additional technical problem for embedded computer applications is that better methods for accessing underlying hardware (low level I/O) from HOLs are probably required before dependence on machine language can be totally eliminated. This language provision will most likely require advances in software technology, even beyond the provisions included in the current Ada effort. Finally, one respondent wrote that a problem was "acceptance: standards follow, not lead, user acceptance." Although subscription to this "axiom" is not realistic, the point is that the Department of the Navy will have to provide the leadership for industry, and Navy program managers will have to encourage use of HOLs in system development.

Any move away from the current acquisition policy of standardization must not compromise the ability of the Navy to fulfill its mission. Current military software systems depend heavily on the machine language architectures of current standard computers. Estimates are that over half of the existing software is written in machine language for the UYK-7, GYK-12, AYK-14, UYK-19, and UYK-20 architectures. It is for this reason, as well as for minimization of life-cycle costs, that an accreditation policy must allow for capture of existing software in the near term. Until the software base of the Navy can

be captured by a HOL commonality, the accreditation criteria must include standardization on these ISAs. However, HOL/ISA standardization is the ultimate objective.

It is important that requirements not be placed on detailed instruction timings. With regard to computer specification, Timmreck writes "one is not (or should not be) interested in nanosecond add-times, multiprocessing, cache memories, microprogramming, bulk core, and other features of modern computers for their own sake, but only insofar as their presence contributes to more economical processing of the workload. ... these and other sophisticated features are so different in different machines that they defy direct comparison."[4] In order to ease technology insertion and enhance competition, machine designers should have the freedom to trade off architectural elements. The subject of the definition of performance ranges is addressed in the next section. Part of the accreditation process must necessarily be a set of ISA verification programs and procedures to validate the compliance of a particular piece of hardware to the ISA standards. These programs should be quite similar to the diagnostics commonly provided by computer manufacturers to determine and isolate hardware problems. Industry response to the questionnaire indicated that relaxation of detailed timing specifications was a necessary, but not sufficient, step in allowing computer designers the flexibility to insert new technology. The necessity to capture existing software tools is well understood, as is the need for public access to verification mechanisms.

The ISA requirement should be an absolute one with subsetting and supersetting within an accredidation cycle forbidden. The reason for no subsetting is that it obviates the capture of existing software. The reason for no supersetting is to prevent noncontrolled proliferation of similar but

different architectures. It is certain that if enhancements to an ISA exist, use of those enhancements will exist, thus nullifying the transportability of accredited systems across programs. Thus, the ISA standard will be updated in a controlled way (presumably supersetting the old standard), and each new standard strictly enforced in the accreditation cycle.

Industrial firms were questioned about the desirability of a prohibition on subsetting and supersetting. The overwhelming response was that a complete prohibition of these techniques was not appropriate. It is understood that some flexibility may be required; however, if supersetting and subsetting are permitted their use must be carefully controlled or transportability will be lost.

3.1.2 Performance Factors. A basic idea behind the accreditation concept is that computer manufacturers will compete their machines against a set of accreditation criteria to win a place on an accredited list. This concept is different from the current concept where competition is against requirements for a specific application. Given the broad spectrum of computer types and capabilities available in the marketplace and the wide range of military embedded computer applications, it is necessary to establish criteria and procedures for classifying computers into performance categories. There are many techniques available for the quantitative measurement of the computational performance requirements of a particular problem or the capability of a particular computer. Some of the more common techniques are:

- o Benchmark program performance
- o Kernal evaluation
- o Instruction mixes
- o Instruction cycle time
- o Memory cycle time

Some of these techniques (kernal evaluation, for example) provide considerably more accurate quantitative evaluation, but require a more comprehensive application analysis than simpler techniques like the comparison of memory or instruction cycle times. For purposes of accreditation, the computational requirement is presumed to be defined by the speed required to execute a specific problem and the identifiable instruction mix peculiar to an application class.

Results of a study performed for the Army's MCF Program (Navy applications were also included) indicate that Navy embedded computer applications can be grouped into five definable classes: command and control, communications switch, number processor, data base management, and character processing.[5]

Command and control designates the operation of command and control functions through subordinate terminals, devices, or computers. The processing requirements normally involve time-sensitive, low-to-moderate arithmetic processing of limited precision, and substantial I/O manipulation.

Communications switch applications employ computers for message, circuit and/or packet switching. Typical installations are characterized by moderate real-time constraints dictated primarily by user response requirements. Arithmetic processing requirements are low, while byte or character handling requirements are moderate. I/O operations are substantial and are usually terminal or support peripheral related.

Number processors require moderate-to-high performance arithmetic processing, often involving multiprecision or floating point. They also are usually required to meet some real-time requirements, such as in navigation or guidance tasks.

Data base management requires a computer which performs control and manipulation of sizable data bases. High reliance on direct access mass

storage is characteristic of such systems.

Character processing applications require substantial alphanumeric data processing. Byte operating capability is an important parameter of such systems.

A result of establishing application classes is the ability to derive a set of performance matrices of requirements relating performance and applications of computing machines which are candidates for accreditation. An example of such a matrix is shown in Figure 2.

Figure 2

PERFORMANCE BY APPLICATION CLASS MATRIX

	Application Class				
PERFORMANCE LEVEL (KOPs)	C^2	Comm. Switch	Number Processor	DBM	Character Processor
HIGH					
MEDIUM					
LOW					

It is envisioned that each cell in the performance matrix would contain at least two accredited machines, in order to maintain a competitive environment. A given machine will likely qualify in more than one application class.

Each candidate machine would be evaluated to find where it would fit into the matrix. This evaluation should be conducted in one of two ways: through use of benchmarks or by establishing equations reflecting the characteristic instruction mixes for each application class. The benchmark for a given application class might be an existing operational program, or it could be a specifically tailored evaluation routine. As the benchmark routine is run on the candidate machine, a performance monitor would be used to measure, for example, throughput, and that computer's place in the matrix is determined. A fundamental drawback to the use of the benchmark evaluation technique is that the candidate machine must actually exist, thus making it difficult and expensive for a manufacturer to attempt an optimization for one specific application area and performance level.

This drawback is not present in the use of instruction mix equations as an evaluation tool. If the application class instruction mix equation is sufficiently detailed, a computer manufacturer can determine the proper spot in the performance matrix simply by knowing instruction cycle times for his machine and exercising the equation. This process can be accomplished without actually assembling a working machine and thus allows the building of machines optimized for a specific application class.

It must be pointed out that great care would have to be taken in evaluating candidate machines with the benchmark technique in the near-term operation of the accreditation program because of the proliferation of ISAs currently in existence. Particularly when using an existing operational program as the benchmark, a candidate machine could be unfairly evaluated

simply on the basis of ISA incompatibility between machine and program. Specifically tailored evaluation routines would ease this near-term problem, as would the use of instruction mix equations.

The problem of noncompatible ISAs will be removed in the long-term operation of the proposed accreditation program, because one of the target goals of this program is an HOL/ISA standardization. Assuming that the proposed application classes and evaluation procedures can be established, there will be a relaxation of the requirements that specific instructions meet fixed timing thresholds. Computer designers will have some flexibility in trading off the speed of some instructions for others, maintaining an instruction mix bandwidth, thereby aiding the designer in producing a machine that would execute in a targeted performance range.

Industry response in this area was mixed, but at least two points seemed to be generally agreed upon. It is felt that definition of performance is worthwhile and that some type of evaluation routine, such as a benchmark, must be available in the public domain.

3.1.3 Level of Hardware Standardization for Accreditation. One of the topics that always creates a controversy when considering a standardization strategy is the level at which hardware standardization should occur. The alternatives discussed here are card, module, or box level.

Most Life-Cycle Cost (LCC) models indicate that maintenance and sparing should occur at the card or module level, but that is not the issue here. The subject of level of standardization for maintenance and sparing is discussed in Section 4.0 of this report. What is being considered is the hardware standardization level to which accreditation criteria should apply. With a box level standardization, functionality of a complete computer would be

specified. The number and kinds of subunits (modules) contained in that box would only be germane to the LCC evaluation of the box from a logistics, sparing, maintenance, and training point of view. Module level standardization would require that modules acquired from different sources be plug compatible in some sense, with interchangeability within a single box. Standardization at the module level seems attractive at first glance, since procurement would take place at potentially the same level as sparing. Competition could take place at this level, and the logistics problem is greatly simplified. Examples of attempts at standardization at the module level include NECS, Standard Electronic Modules (SEM), and the initial MCF concept.

In order to standardize at the module level, form/fit/function constraints are required at the module level. In addition, in order that the modules will be plug compatible, a standard bus definition is required. Proponents of module level standardization point to the existence in the commercial marketplace of second-source suppliers of memory and peripherals for existing machines. Given the existence of emulation as today's technological answer to upward compatibility and ISA implementation, module standardization may be easily extended to require compatibility across different ISAs. Examples are the NECS and MCF programs.

Proponents of box level standardization are quick to point out that module compatibility across multiple ISAs has yet to be proven. The argument is that to achieve efficiencies required by existing ISAs, the designer must not be forced into conforming to a fixed bus standard or, for that matter, any other forced partitioning of the components of the system. The assertion is that module level standardization will stifle any large advances obtainable through technology improvement, as well as remove incentives for industry to participate in such developments. In looking at computer architecture

implementations in industry, it is not difficult to find large variations in bus architectures even between different performance members of the same ISA family, e.g., the PDP-11 family.

The computer industry questionnaire received a unanimous response in favor of standardization at the box level for accreditation. Respondents listed several benefits accruing as a result of box level standardization:

- o simplification of configuration management and control
- o elimination of expenses associated with computer integration
- o single supplier responsibility for operations and support

The need for functional and or interconnect standards was pointed out.

Industrial firms were asked if standardizing on a bus would allow enough design freedom to incorporate technology advances in new implementations of a standard ISA. The range of responses suggested a general feeling of discomfort with bus standardization. Respondents replied that bus saturation often defines a limit on throughput and that bus standardization is too sensitive to technology advances to be frozen over a long period of time.

Once a new computer is accredited, it might be possible to recompete modules that make up that member on a form/fit/function basis. One of the consequences of such a policy is that profit incentives are reduced or removed for industry to participate in the R&D required to design new accredited computers. In response to a question regarding the desirability of separately competing modules that make up a box, the computer manufacturers indicated disfavor. The problem areas associated with multiple suppliers of modules are contrary to the benefits described above for box level standardization:

- o configuration management and control
- o need for expenditures to ensure proper integration
- o assignment of responsibility for poor performance by a module designed to specs

A suggested alternative is to pursue accreditation of a design first, then consider second-sourcing of modules by either private or public competition, depending upon who owns the design rights. In this situation, module competition would be a result and not a preexisting condition of accreditation.

It is our view that the target level of standardization for accreditation criteria should occur at the box level. Thus designers will have complete flexibility in partitioning processing functions in order to obtain desired performance. This does not preclude a vendor from offering a set of computers which uses standard modules among them, thus decreasing government LCC and making their use more attractive. The accreditation program should avoid specifications in accordance with known existing technology which may change in the near future.

A MIL Standard, such as 1397, should be used as the standard for computer interconnections. Standard interfaces to external sensors should be defined to allow interchangeable deployment of accredited computers.

3.1.4 Life-Cycle Cost. Life-cycle cost (LCC) will likely become the most important criterion in determining which computers are accredited. Most of the criteria discussed to this point are on a pass/fail nature with respect to attaining accreditation, whereas LCC will rank the various candidates in competition for entry to the accreditation list.

Life-cycle costs include such things as recurring and nonrecurring costs

of the production investment, operating and support costs, and research and development costs. In order to enable discrimination between competing candidate machines, an accurate but versatile LCC model must be developed. The LCC model must be accurate so that incremental effects on LCC with respect to variations in performance characteristics, machine definitions, reliability/maintainability, technology, etc., can be easily and independently observed. The LCC model also must be sufficiently versatile to allow it to reflect the real world today and into the future so that the best of competing alternatives may be selected without bias from inaccurate or incorrect cost estimating relationships. LCC models have been developed for TRI-TAC and other current programs; however, no model currently exists which is capable of performing the required task for accreditation purposes.

The computer manufacturers were asked if current life-cycle cost models were adequate for purposes of accreditation. The responses indicated general dissatisfaction with available LCC models for a variety of reasons. Among them were a perceived need for updated models to addresss new technology designs and maintenance methods; a desire for more detail in the areas of man hours to make repairs, percentage of repairs made at each maintenance level, number of spares available, MTTR and MTBF; and concern over methods to verify the credibility of model input data.

The major factors that contribute to the life-cycle cost of a logistics support system for embedded computers are:

- o Repair costs
- o Inventory costs
- o Specifications and drawings costs
- o Transportation costs
- o Training costs

- o Test and diagnostic equipment costs
- o Technical manuals costs
- o Personnel costs
- o Facilities costs
- o Repair parts costs

The life-cycle cost for a logistic system is simply the sum of these individual components. To create a useful cost model, it is necessary to break these components down even further and estimate their costs individually. Several of the factors are constant costs that are incurred one time for a particular system. Some costs depend on the quantity of computers purchased, and yet others depend on the number of people required to service the computers. An example of the level of detail of necessary in the envisioned LCC model is shown in the next paragraph.

Repair costs are *directly proportional to the number of failures* processed by the repair facility. Costs related to repair such as the test equipment and the personnel are treated in the other factors. The costs for repair are modeled by the equation:

$$\text{Repair costs} = R.N.LT/MTBF$$

where

R is the cost to repair one item

N is the number of items in use

LT is the life-time of the embedded computer system

MTBF is the mean-time between failures of an item in active use.

The factor N.LT/MTBF is simply the number of failures that occur during the lifetime of a system.

This formula is a fairly crude but useful measure of the repair costs.

There are second order cost factors that can be incorporated as indicated below, but these may simply add unnecessary detail to the cost model and detract from its essential purpose of providing an easy means of estimated logistics costs.

The second order effects account for additional failures (failures of inactive systems while out of service) and varying costs for repair. The additional failures are:

$$\text{Repair costs for inactive equipment} = R.N_I.LT/MTBF_I$$

where

R is the repair cost for a failure

N_I is the number of inactive systems in spares and warehouse

LT is the lifetime of the embedded-computer system

$MTBF_I$ is the mean-time between failures for an inactive system (the shelflife of the system)

Although failures of inactive systems are rare, they do occur because of such problems as pins and contacts making poor connection, improper storage environment, mishandling of cabling and interconnections, and similar other problems. Thus to estimate the costs for various kinds of failures, it is necesssary to breakdown the repair costs by failure category and sum them over the failure types using the equations above.

To develop the desired LCC model, it will be necessary to derive detailed relationships for all of the major factors that contribute to life-cycle cost, as in the foregoing example for repair costs.

3.2 Number of Computers on Accredited List

In previous paragraphs a scheme was described for evaluating candidate computers and placing them into a performance matrix for purposes of accreditation. Under the current Navy acquisition policy of standardizing on one computer in a performance range, the suggested performance matrix would have only one very broad, generic application class with one entry at each performance level, as shown in Figure 3.

Figure 3

PERFORMANCE MATRIX UNDER CURRENT POLICY

	Application Class
HIGH	X
MEDIUM	X
LOW	X

In terms of the accreditation concept, the performance matrix in Figure 3 shows three accreditation lists (corresponding to the spaces in the matrix), with one computer entry in each list (corresponding to the "x" in the matrix space). Under the proposed accreditation program, the performance matrix would be expanded to include five application classes and and the same three performance levels, as shown in Figure 4.

Figure 4

PERFORMANCE MATRIX UNDER PROPOSED ACCREDITATION PROGRAM

Application Class

PERFORMANCE LEVEL (KOPs)	C^2	Comm. Switch	Number Processor	DBM	Character Processor
HIGH	X_1 X_2 .... X_n				
MEDIUM					
LOW					

The result would be fifteen accreditation lists and each list would contain some number of accredited machines. An attempt was made to determine how many accredited machines or slots are appropriate in the accreditation lists.

Under the current acquisition policy, the number of available slots in each list is one. Problems associated with the current policy of standardization, as described in the introduction to this report, would indicate that one slot per list is not sufficient. The stated reason for the current Navy policy of standardization on one computer per performance range is hardware maintainability. The need to operate, maintain, and spare embedded computers on ships at sea has had a tendency to limit the proliferation of computer types in use. It would seem logical then to examine the effect upon logistic costs of having more than one computer type in use.

In 1978 and 1979, a working group consisting of representatives of all three services cooperated in a qualitative analysis of logistics life-cycle

costs incurred in the support of military embedded computer systems.[6] The objective of this effort was to determine the principal cost factors associated with two repair concepts: warranty/contractor and in-service. The Army funded another study in 1979 which attempted to quantify the cost of spare computer components as a function of logistic support concepts and to model the effects of multiple suppliers on spares costs.[7] The pertinent factors that contribute to logistics costs were identified as the following:

- o Contractor support costs
- o Inventory (pipeline and float)
- o Transportation
- o Repair parts
- o Personnel, training and facilities
- o Specifications, documentation, technical manuals, test and diagnostic equipment

It was determined that the major factors that contribute to life-cycle costs for contractor repair are:

- o Contractor support
- o Inventory
- o Specifications and documentation
- o Transportation

For in-service repair, the major cost factors are:

- o Personnel, training and facilities
- o Specifications and documentation
- o Inventory
- o Repair parts
- o Transportation

After analyzing the major cost factors for both repair concepts, the study reported that the costs associated with specifications and documentation will likely have the strongest effect when considering multiple vendors. These are direct costs for both contractor and in-service strategies, and they

are also reflected indirectly in the contractor charges for contractor repair where they cover in-house costs for items that are not deliverables.

A secondary area in which costs depend on the number of suppliers is the cost of spares. Our analysis shows that these costs grow slowly with the number of suppliers, which indicates that the multiple supplier cost burden is more likely to be felt in terms of documentation and specification. This topic is discussed further in Section 4.0.

The intent of this analysis of life-cycle cost factors was to reveal some critical point in the relationship between the number of suppliers and the resultant logistics costs. If such a point could be found, it was to have been used as an indication as to an upper limit on the number of slots available on each accreditation list. However, because of limitations of available cost data, the curves that were generated did not exhibit any identifiable "knees" or break points. It is felt that with more time and resources to devote to data collection and analysis, some meaningful guidelines could be developed. For example, data relevant to the savings accruing from competition would be extremely valuable.

3.3 Accreditation Cycle Length

Accreditation cycle length refers to the periodicity with which the accreditation lists will be opened up to consider new candidate machines for accreditation. Three of the stated goals of an accreditation program for computer acqusition are to stimulate competition, ease technology insertion and shorten the acquisition cycle. As with any optimization problem, there are at least two points of view or approaches to the problem which normally give conflicting results. From the user's standpoint, there is certainly a great interest in shortening the time required for acquisition. This would

result in a better rate of technology infusion and reduce life-cycle costs. However, because of sunk costs involved with maintenance training, documentation, setting up logistic support, etc., the user would prefer a longer, as opposed to shorter, accreditation cycle length. Conversely, from the supplier's standpoint, a shorter accreditation cycle would be preferable because of increased opportunity for sales. As new products are developed they could be more rapidly offered for accreditation with a shorter cycle length.

3.3.1 Factors Affecting Accreditation Cycle Length. There are two factors affecting the accreditation cycle length. First, under the Navy's present acquisition policy of standardization, it takes an average of about seven years to complete the cycle of buying a new computer for embedded applications. On the other hand, computer technology is advancing at such a rapid rate that by the time an acquisition process can be completed, the newly acquired computer is obsolescent with respect to what is then currently available in the marketplace. The rate of advance of computer hardware and software technology should have some bearing on accreditation cycle length. While it may be difficult to evaluate precisely, a review of the pace and trends in computer technology advancement should reveal some quantitative insights which might be a guide to cycle length.

Second, life-cycle cost is an important element in determining accreditation cycle length. Because of the identifiable costs associated with setting up a logistics system to support a new computer acqusition, the user cannot afford to be accrediting new machines at too short an interval. On the other hand, there are costs associated with operating and maintaining depreciating equipment. From a cost standpoint, the key to introducing new

technology is that the future benefits of the new technology must be greater than the present costs of introducing it. To the extent that benefits outweigh costs, it is worthwhile to introduce new technology. However, if gains are small and introduction costs are high, it is better to retain older technology. Finally, the rate of introduction of new technology cannot be too fast, because new systems must be installed and stable for a number of years in order to derive some cost benefit.

3.3.1.1 Computer Technology Advancement Rate. In attempting to survey the advancement of computer technology, one discovers that there are two schools of thought on the subject. One group holds that there have been and will continue to be recognizable, major (revolutionary) advances in technology. The other group maintains that technology advances are purely evolutionary in nature. Optimization of production machines to specific applications causes a sort of evolutionary improvement in technical capability. However, incorporation of new developments in processor and memory chips, peripherals, and I/O architecture and power supplies, for instance, may be regarded as providing more revolutionary improvements in technology. Using this latter viewpoint, our analysis of the literature indicates a three to five year cycle in major changes or updates in computer hardware technology.

The survey of industrial firms indicated that the identifiable periodicity in the introduction of major changes in computer technology ranged from two to four years, and certainly no more than five years.

3.3.1.2 Life-Cycle Cost Considerations. In an effort to investigate the relationship between system life-cycle cost (LCC) and accreditation cycle

length, a cost model was developed that shows the effect of new technology on LCC for embedded computer systems. The model uses a present value of future expenditures, thereby discounting both future costs and savings to reflect the greater value of present monies.

It is assumed that there are three cost components to an embedded computer system that determine its LCC: research and development (R), procurement (P), and annual logistics costs (L). The logistics component includes all annual expenditures such as spares, maintenance personnel, training of maintenance personnel, inventory, transportation, and warehousing costs. A key assumption in the model is that technology improvements occur on a regular basis, thereby increasing some future savings by using the most modern technology possible. To model the effect of time on technology, a technology improvement factor is included. It is assumed that each year one can purchase equal capability in computers for a fraction less than the cost in the previous year. Similarly, logistics costs for new systems will be fractionally lower each year due to addition of built-in-test, higher reliability, smaller size and weight, etc. R&D is modeled at a constant cost in fixed dollars and does not decrease with technology improvement.

To determine the LCC affect upon the accreditation cycle length the point at which logistics savings is maximized must be determined. This expenditure must then be balanced against discounted R&D and procurement expenditures. Generally speaking, discounting and technology improvement factors for the costs of R&D and procurement lower as time increases. The optimum point for inserting new technology will be sometime after the point at which logistics savings are maximum, since this later point in time may achieve lower total cost from lower R&D and procurement costs, in spite of the slightly reduced logistics savings.

The LCC model was exercised for a range of values of discount factor (d), technology improvement factor (t), and year of introduction (k). The discount factor values used were 5%, 10%, and 15%. For most present value calculations, a discount factor of 10% is used in accordance with DoD Instruction 7941.3. However, in recent times there is strong evidence that 10% may be too low over the next two decades.[8] Technology improvement factor values used in the model range from 5% to 25% in increments of 5%. For some aspects of technology the historic trend has been as high as 20% per year, most notably in memory technology. However, not all aspects of device technology have shown this improvement, and there is some doubt that military technology can improve at the rate of 20% per annum in costs over long periods of time. In light of these considerations, it is felt that the values studied should bracket the possible range of such factors over the next several years. An example of one of the model calculations is shown in Figure 5.

On the basis of the model runs over several sets of parameters, there is strong evidence that the accreditation cycle should allow for the introduction of new technology about six to eight years after the initial accreditation of a machine in an application class. The primary observation is that logistics cost savings are maximized in every calculation for periods of time in the four to seven-year range. The actual savings are somewhat less than the logistics savings because of the need to account for R&D and procurement expenditures. Savings in logistics tend to be maximized in the four to seven year time frame; however, total life-cycle costs will be minimized at some point in time slightly later than the maximum logistics savings time in order to reduce the cost of R&D and procurement. A more rigorous presentation of this cost model analysis appears in Appendix C to this report.

Figure 5
LIFE CYCLE COST RELATIONSHIP

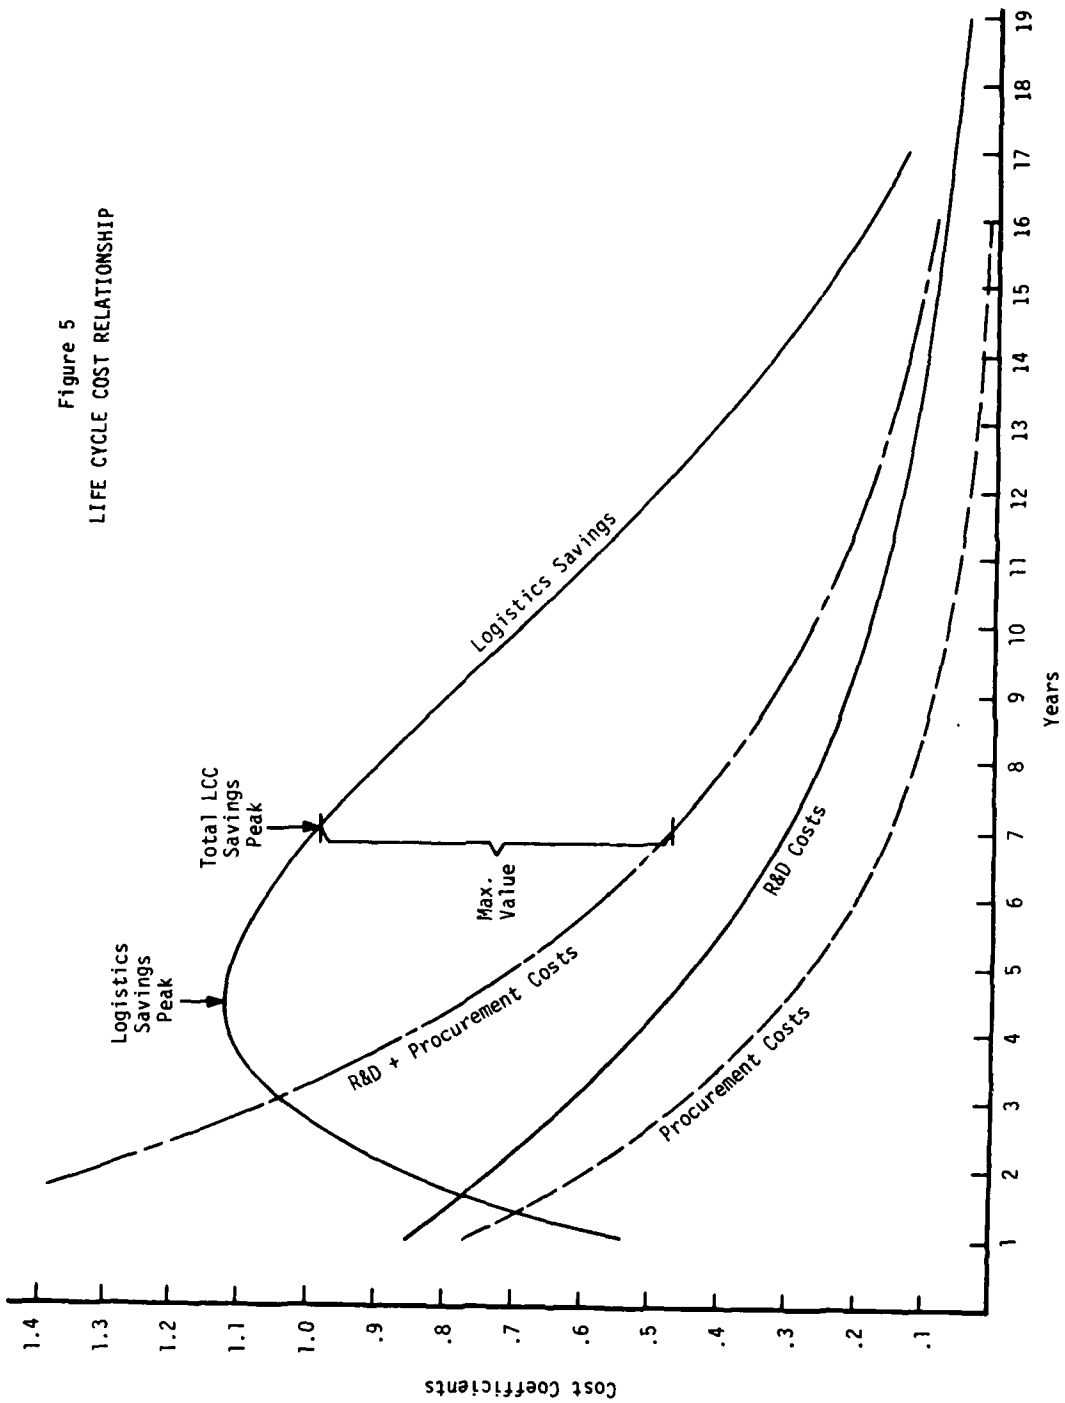

3.4 Criteria for Removal from Accreditation Lists

Among the objectives in designing an accreditation program are stimulation of competition and easing of technology insertion. To assure that the accreditation program will not stagnate to the detriment of obtaining these two objectives, there is a need for specific criteria by which previously accredited computers can be removed from accreditation lists.

It would seem appropriate to allow a manufacturer to withdraw his own product from an accreditation list. However, to ensure maintainability of any of those products then in service, there should be some requirement included in the accreditation criteria for continued logistics support, at least until the end of the current accreditation cycle. Another criterion for removal from an accreditation list might be failure to maintain specified performance standards for that list, or failure to maintain MTBF requirements in service. LCC is another potential removal criterion. At the end of the accreditation cycle, when the accredited lists are opened up for competition, candidate machines which are otherwise qualified would be ranked in order of increasing LCC. Assume there are three machines on the list and two new candidates, for a total of five candidates seeking three available slots on the accreditation list. The three machines with lowest LCC would be included on the new list. With guidelines such as these, competition could keep the number of machines on a given accreditation list within a reasonable bound, even if no fixed upper limit is applied. A manufacturer who is not selling any machines in a specific performance category is unlikely to bear the expense of retaining his equipment on that list, and he will remove his product. At the end of the current accreditation cycle, that manufacturer could once again compete his products for inclusion on the appropriate lists.

4.0 MAINTENANCE CONSIDERATIONS

This section discusses several topics associated with the maintenance and sparing of embedded computers on ships at sea. The increasing use of embedded computers and the additional numbers of types of computers which may result from an accreditation policy require careful consideration if logistic costs are not to overcome the benefits of accreditation.

4.1 Built-in Test

This section discusses the basic concepts of built-in test (BIT), techniques currently available for the implementation of BIT, problems and costs arising from the use of BIT, and the impact of BIT on the maintenance of embedded computer systems.

4.1.1 The Concept of Built-in Test. Built-in test is characterized by the design of computers in such a way that testing is an integral part of design on all levels -- hardware, firmware, and software. On-line monitoring of the internal processes of the computer is provided, thereby increasing observability and offering verification that the computer is indeed working correctly.

When a fault occurs in a system, several things must occur before the fault can be repaired. First the fault must be detected. Without BIT, this often means that some human must become aware that something is wrong, either because things stopped happening or because something happened that should not have. This approach is an expensive way to detect faults. Second, the source of the fault must be isolated. Traditionally, this involves systems personnel and repair technicians in tracing the state of the system at

the time of the fault and locating it with sophisticated test equipment. This approach can also be expensive, not only in labor costs, but in downtime, especially if availabilty of the system is a critical factor. Once the fault has been detected and isolated, the identity of the faulty component must be communicated to someone or something which can act on that information to effect a recovery or repair. BIT is designed to automatically provide these three operations of detection, isolation, and communication.

BIT does not itself make parts more reliable, (i.e., less likely to fail) rather it provides a basis for response to failure. This response in turn can lead to enhanced reliability of the system as a whole. BIT may be used as the basis for a fault-tolerant system which might achieve greater reliability by responding to faults in a variety of methods, such as reconfiguring the system using redundant modules. It may be used to improve the maintainability of a system through expediting repairs or it may be used to reduce the need for and ease the load on off-line automatic repair equipment.

Effective and efficient design of BIT requires an understanding of system requirements and objectives and a knowledge of relevant fault populations and their associated rates of failure. Faults may be divided into two kinds, permanent and intermittent. Permanent faults, as the word implies, are those which remain faulty and may be caught by periodic testing. Intermittent faults may appear and disappear; for example, an open circuit due to a loose connection, or a change in a VHSIC component due to the impact of an alpha particle. According to a study performed for the Naval Electronics System Command by Clary and Sarone, this type of fault accounts for as much as 90% of all faults in some systems and may account for more than 90% of all maintenance expense due to the greater difficulty in detecting and isolating

them.[9] Concurrent testing techniques are generally required to handle intermittent faults effectively.

Modular composition of a system according to function can greatly enhance detection, isolation, and communication of faults. Combined with BIT, modularization can also lead to easier and faster repair through the replacement of faulty modules identified through BIT.

4.1.2 Built-in Test Techniques. Approaches to BIT can be divided into concurrent and nonconcurrent techniques. Concurrent testing is performed throughout the same time period as the execution of the primary processes in the system. Nonconcurrent testing operates in between the execution of the primary processes in what would otherwise be idle time.

Nonconcurrent techniques are useful primarily for the detection of permanent faults and may be implemented through software and firmware. Software tests generally require no additional hardware and may be callable not only from the operating system but also the user software. Firmware may be used to generate test patterns and test reference data, or it may be used for microdiagnostics in accessing individual gates, paths, and circuits.

Concurrent BIT techniques may be especially appropriate for systems operating in real time, as is generally the case with embedded computer systems. They are also the only effective way to catch intermittent faults. These techniques are generally based on the principle of redundancy and add to the cost of hardware.

Information redundancy, Hamming codes and constant ratio codes are examples of software-related diagnostic techniques available as applications of BIT.

In hardware, circuits may be designed to check themselves at the gate

level. The principle of redundancy may be applied through the replication of parts of the system. Modules such as processors may be replicated, or even the entire computer may be replicated. For BIT, dual redundancy (having two of the same part) is generally sufficient, since this allows for the comparison of results (voting). Higher order redundancy may be used for fault-tolerant systems.

In modularized computers, software and firmware may be used to isolate functional paths leading to recognizable replacement modules. Then microdiagnostics can be used to access individual gates, paths and circuits to indicate the component to be replaced.

To questions concerning BIT on the industrial questionnaire, all respondents indicated that they used BIT at both the chassis and module levels, forty percent used BIT at the board level, and twenty percent used BIT at the integrated circuit level. Those questioned were generally in favor of increasing their use of BIT, but some indicated that they would do so only if user requirements forced the investment.

4.1.3 Problems and Costs Associated with Built-in Test. In general, one would expect the inclusion of built-in test to increase design, development and production costs because of the need for additional hardware and software. A study of the cost-effectiveness of self-checking computer design performed by IBM scientists on the S/360 computer, concluded that sixty-five to eighty percent of faults were checked by a thirty-five percent increase in hardware over an unchecked S/360 computer. This degree of checking was accomplished without degrading the performance or speed of the machine.[10]

With the expected benefits of the DoD's VLSI and VHSIC programs and the ever decreasing costs of hardware, the problems and costs of BIT do not appear

particularly significant when viewed in light of the life cycle support savings which could result. (See section 4.1.4 for implications of maintenance savings accruing to BIT).

Of course, BIT must be carefully developed and clean, well documented interfaces established between the testing and operating system and the user. Furthermore, the communication of fault information between modules of a modular system, where if modules were supplied by multiple vendors, would have to be standardized.

4.1.4 Built-in Test and Maintenance. To maintain a computer in running condition, faults must be detected and isolated and the faulty element must be either repaired or replaced. As seen in the discussion above, built-in test provides for automatic fault detection and isolation. Automatic detection of faults may eliminate the need to have someone always on hand to watch for faults. A few technicians in a central location might be able to monitor and respond to the needs of a number of systems as opposed to requiring at least one technician per system. This approach might result in the reduction of the total number of technicians required. Automatic detection may be more accurate and result in fewer false alarms. Additionally, automatic detection is often faster, so that less time is spent in improper functioning before detection of the fault. Thus, automatic detection may result in greater availability of the system, and it might permit repair before the fault begins to spread to other areas. Automatic and rapid fault detection also serves to increase confidence that the system is functioning properly.

Automatic fault isolation can cut the time for repair significantly through a great reduction in the time required to isolate the fault. It

can also result in fewer "trial and error" repairs where replacement parts are inserted until the problem goes away. This kind of repair may be excessively costly in both parts and time. Less external diagnostic support equipment is needed when the system is capable of isolating its own faults. Self isolation of faults can also provide more accurate information on fault populations within the system, thus yielding a better basis for estimating spare parts needed. Fewer technicians may be needed overall, as the need for diagnosing faults is reduced. Perhaps the most significant advantage of the automatic isolation of faults is the reduction of the skill level required of technicians since they would no longer need to know how to do fault isolation themselves. With a reduced skill level, technician trainees could be drawn from a broader base within the civilian population. Training would also become less expensive, since less equipment would be necessary and less time would be spent in training. Estimates of the current costs for training a DS3 are about $23,000 for billet and $5,000 to $7,000 for educational expenses, over about 70 weeks. A shorter training period would mean more time in the field after training and would permit a faster response in supplying technicians as demand increases. Furthermore it is possible that these less skilled technicians would be under less pressure to leave the Navy for higher paying jobs in industry and might be more readily induced to reenlist.

4.1.5 Significance of BIT to Accreditation. In the studies conducted by Stone and Kohler, [11, 12] it was revealed that for the in-service repair concept (repairs provided entirely by Navy technicians at the field and depot levels), the impact of multiple suppliers of computers on personnel costs associated with logistics and maintenance is strongly dependent upon the

effectiveness of built-in test. If built-in test features are effective, then personnel costs may be largely independent of the number of suppliers. Otherwise, personnel cost can become proportional to the number of suppliers, resulting in a linearly increasing cost for logistics as the number of suppliers increases. Thus, the requirement to incorporate effective built-in test features in new generation embedded computers is vital to the success of accreditation as an acquisition policy.

4.2 Mean Time Between Failure

Mean Time Between Failure (MTBF) is an elusive, hard to define, difficult to quantify measure of system reliability. MTBF is very sensitive to the environment in which a computer operates and also to the operational schedule. Experience has shown that a given computer, regardless of its stated MTBF figure, will operate for a longer period between failures if it is kept cool and operated continuously, as opposed to cycling the machine on and off and operating it in a high temperature environment. In spite of the problems associated with uncertainties about MTBF, it is a useful concept. Given a constant environment and operating schedule, a more reliable computer (longer MTBF) will result in lower logistics cost than a less reliable one. The tradeoff for longer MTBF is in higher procurement costs, because it generally costs more money to ensure greater reliability.

Against this simplistic background one might be tempted to simply pay more money for longer MTBF in an attempt to eliminate the need for maintenance technicians and spare parts on ships. According to OPNAVINST 4441.12A, Navy ships are stocked for a maximum endurance of only 90 days, which translates to a little over 2000 hrs. It does not seem unreasonable to assume that

technology could provide a computer which would operate for 2000 hours without a failure. However, the concept of MTBF refers to operation in a fairly benign environment and no operating guarantees can be provided for a warship going into combat. Thus, it will always be necessary to have some organic maintenance capability and spare parts on Navy ships. Even so, it should be possible to predict a value of MTBF which would be affordable and result in a reduction in numbers of maintenance technicians required to be aboard ship.

The Naval Underwater Systems Center (NUSC) was requested to provide real-world observed data on MTBF for embedded computers currently in operational use. NUSC reported the following figures, as of November 1979.

Computer	MTBF
AN/UYK-7 single bay	1800 hours
AN/UYK-7 three bay	1200 hours
AN/UYK-20	2000-2400 hours

In April 1980, the Office of the Assistant Secretary of the Navy for Research, Engineering and Systems reported that observed MTBF for the AN/UYK-20 was in the 6000-7000 hour range. No data is yet available on observed MTBF for the AYK-14, but 2000 hours MTBF has been specified for that machine. Available performance data indicate that current technology can provide a MTBF for embedded computers which equates to the 90-day planned endurance for Navy ships. However, NUSC also reported that an availability rate of 85% on non-critical functions and 100% on critical functions is desired. These levels of desired availability have the effect of increasing the required MTBF to provide that level of availability over the 90-day ship

operational period. If the distribution of failures can be determined, from analysis of real world data or scientific deduction, the level of MTBF needed to meet specified availability over the 90-day operating period can be computed.

4.3 Software Maintenance

One of the additional tasks facing shipboard maintenance personnel is the installation and maintenance of software in embedded computer systems. If this task could be handled in some other way, it would contribute to a reduction in training requirements for shipboard maintenance technicians and also in the required basic intelligence level of technician trainees.

It is proposed that software programs be thoroughly validated or debugged in the shore establishment, possibly by civilian technicians. A designated traveling team of specially trained military technicians could carry the software into the fleet to install it and check out the embedded computer system on the ships. A program would have to be established whereby less skilled technicians aboard ship could report software "glitches" to the shore establishment for attention.

4.4 Level of Maintenance and Sparing Standardization

Section 3.1.3 of this report discussed the appropriate level of hardware standardization for purposes of accrediting computers. It was concluded that to enable the greatest designer freedom in incorporating new technology, standardization at the box level is appropriate. There are different considerations, however, when considering maintenance and sparing aboard a ship.

BIT has the potential to allow a relatively low skilled maintenance

from the ship's onboard spare parts supply. BIT has the capability to identify a faulty box, module or card, thus BIT does not limit the level at which organizational maintenance may occur. Two factors which do have a deciding effect upon level of maintenance are the logistics costs associated with maintaining the inventory of spare units and the limitations in storage space on Navy ships, particularly on submarines.

Studies have been conducted to investigate the relationship between the logistics costs of repair units and unit failure rate.[13] Failure rate is described as the ratio of supply cycle time to MTBF for faulty replaceable units. Supply cycle time is the period required to remove a failed unit from a deployed system, ship the unit to its point of repair, repair it, and then return that unit to the spares inventory of a deployed system. For a submarine, the supply cycle time includes the time of a typical patrol plus the time for cycling a failed unit from a port, through the repair process, and back to the port. The failure rate provides the basis for a determination of the average level of spares required to be assured that sufficient spare units will be on hand to continue operations until failed units can be repaired and returned to inventory. For example, if the failure rate is 2.1 per cycle, the spares inventory should have at least three spares on hand (fractional numbers of spares must be rounded up). For practical reasons, this number of spares has to be increased to account for short periods in which the instantaneous failure rate may be higher than the statistical average failure rate. If the spares are stocked at this computed level, the quantity on hand should be sufficient to cover the failure of the item and all subsequent failures that occur before the original item is returned to the spares inventory.

Figure 6 is an example of spares inventory cost as a function of failure

rate, calculated from data on Navy systems. The absolute value of the dollars shown is not as important for this example as is the relationship between the spares inventory costs of modules versus boxes. In this example of Navy system data, the boxes cost about $70,000 each, compared to a module cost of about $7,000 each. These figures are very representative. Realistic failure rates are in the region from .05 to 1.0. Above failure rates of 1.0, systems tend to be viewed as unreliable. Note that the cost for box spares exceeds the cost for module spares at every failure rate above .05. To look at a worse case example, assume a submarine is operating on a 90-day patrol (2,160 hours). By definition the absolute shortest supply cycle time is 2,160 hours plus repair time. Accordingly, assume for illustrative purposes a supply cycle time of 2,200 hours. NUSC has reported that the longest specified MTBF for a typical Navy embedded computer is 2,000 hours. These two numbers yield a failure rate of 1.1. Using the curves of Figure 6, the resulting cost difference between box sparing and module sparing is shown to be about five-fold. This cost difference, plus the severe limits on storage space in a submarine, clearly favors sparing and maintenance at the module level (in this context a module may consist of as little as one card). The supply cycle time for surface vessels is likely to be far shorter than for submarines and storage space restrictions are not as severe. Assuming an MTBF of 2,000 hours and a failure rate of .05, supply cycle time would be 110 hours, a little over four days. This figure is probably very optimistic for a ship operating at sea. Therefore, the relationship of the cost curves in Figure 6 still favor sparing and maintenance at the module level.

Figure 6

SPARES INVENTORY COSTS AS A FUNCTION OF FAILURE RATE

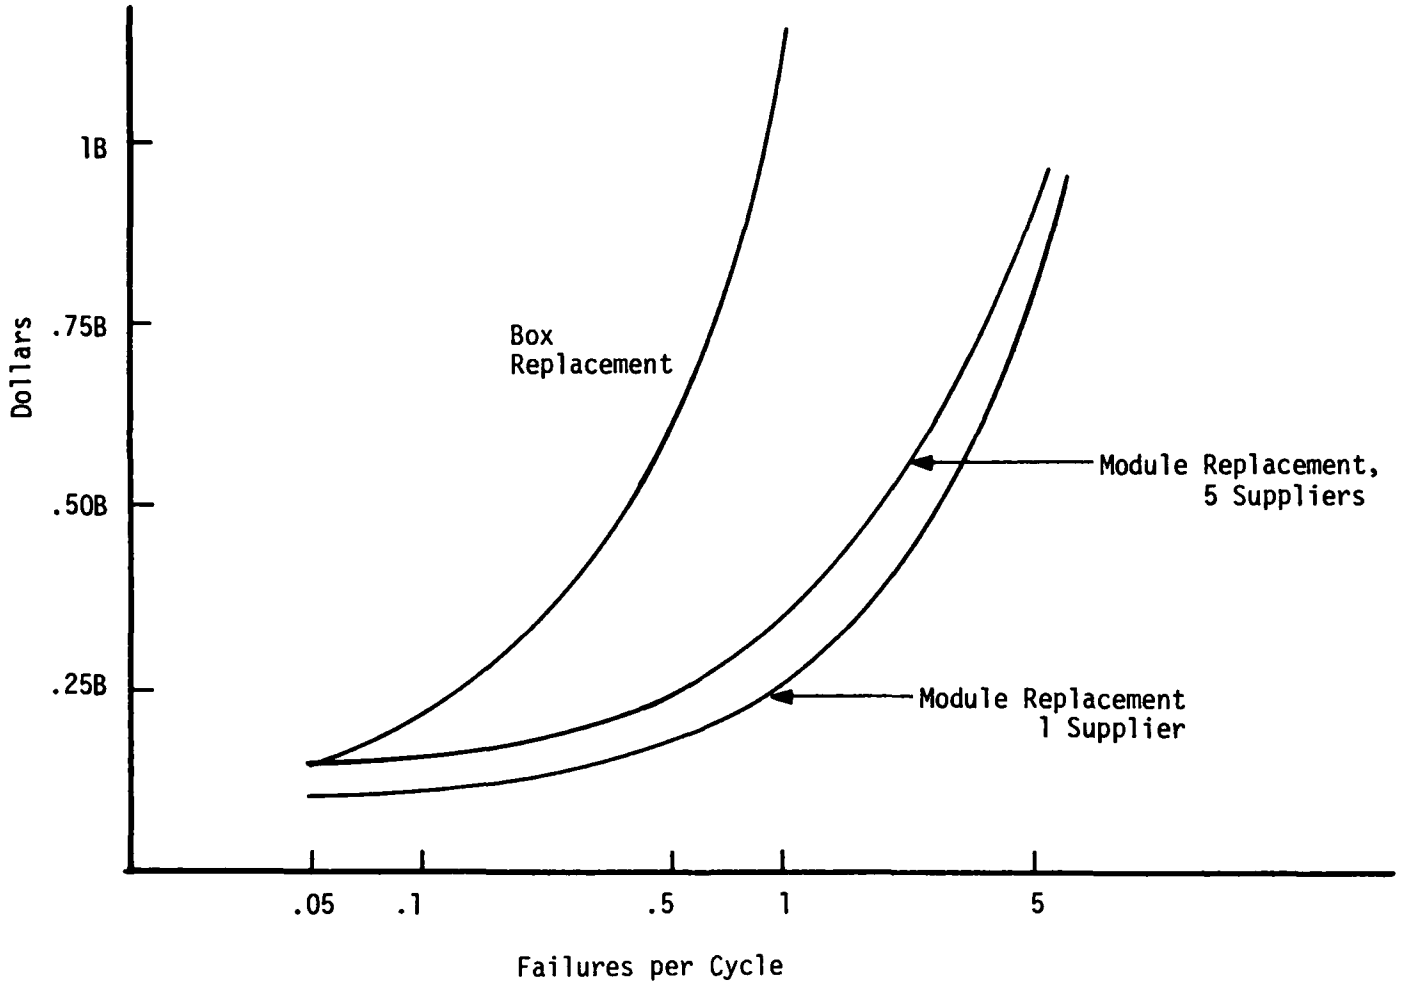

4.5 Box and Module Standardization

It was concluded in Section 3.1.3 that accreditation criteria be met at the box level. In Section 4.4, it was concluded that, for maintenance and sparing, the module is the key level of standardization. This approach appears to have several benefits. First, by accrediting at the box level, opportunities for techology insertion and manufacturers competition are enhanced as a system approach may be taken toward design improvement. Second by sparing at the module level, the cost of sparing and logistics is lower, and the level to which BIT should be able to detect and isolate faults is tacitly set. Third, as technology results in more gates per chip and therefore, more functions per square inch, the number of modules per box in the logistics pipeline will decrease in spite of the increase in box types resulting from the accreditation approach. One may expect that within the decade, a module could equate with a box. As an example of this trend, it is now possible to acquire a computer containing approximately 20 cards which has performance comparable to that of a 1-bay AN/UYK-7 computer containing 800 cards.

5.0 TRANSITION PLAN

This section contains a discussion of factors affecting the timing of implementation of an accreditation program, a description of the recommended accreditation program along with a list of the recommended accreditation criteria, a milestone chart which correlates the criteria with an implementation schedule, and a discussion of the need for government leadership.

5.1 Factors Affecting Timing of Implementation

The Navy has a considerable investment in operational software in the fleet today. It is impossible from economic and operational standpoints to make an instantaneous transition to the target accreditation program. Even if the Navy were willing to pay the price, in dollars and in the threat to our national security, technology would not support an immediate change. A number of factors impinge on the rate at which the accreditation program can mature.

The procurement program currently under way for the acquisition of AN/UYK-43 and AN/UYK-44 computers forms a natural setting for the introduction of an accreditation program in lieu of allowing another sole-source situation to develop. If initial accreditation criteria were selected based upon the acquisition criteria for these two types of machines, the problems of transition from current policy to accreditation could be minimized.

The requirement to incorporate effective built-in test features in new generation embedded computers is a vital one, in order to keep logistics costs from becoming a significant burden under accreditation. While BIT will not likely perform to the desired level today, estimates provided by questionnaire respondents regarding the time to develop and implement new technology

indicate that the necessary capability could be available in approximately five years. This assumes, of course, that the industry is notified well in advance of the requirement.

A key requirement in the target accreditation program is standardization on an HOL/ISA. The proposed HOL standard is Ada. While the specifications of the Ada language will be complete by June of this year, work has not yet begun on an Ada compiler. Both the Army and the Air Force are sponsoring development of Ada compilers, with the Army project slightly ahead of the Air Force. Assuming no further major obstructions to the Army Ada compiler development program, the product of that program should be completed and validated by approximately May 1983. It will not be until this date that an Ada compiler will be available in the public domain. The industry questionnaire asked for an estimate of the number of years in the future when an HOL embedded computer standard might be appropriate. Responses ranged from immediately to the mid-1980's, as far as establishing a standard. However, the answers were caveated by requiring the existence of an Ada compiler in the public domain.

Finally, the requirement to develop a program which will reduce costs over time is of major importance. Parametric analysis of the life-cycle benefits of new technology indicate a maximum savings opportunity every five to eight years after a new technology introduction.

Taking all of these factors into account, it appears that an accreditation cycle length of five years is appropriate. This period provides a reasonable compromise between the rate of technology advancement and the time requried to attain acceptable life cycle cost savings.

Accordingly, criteria for accreditation may be established at five year intervals beginning with interim criteria appropriate to the initiation of the

accreditation program and progressing through mid-term criteria to the target criteria appropriate to a mature accreditation program.

5.2 Accreditation Criteria

Sections Three and Four of this report discussed accreditation factors and maintenance schemes associated with embedded computers. These topics were discussed separately for ease of presentation; however, within the context of an accreditation program for the acquisition of computers, the two subject areas are mutual contributors to the list of accreditation criteria. The collection of recommended accreditation criteria is presented below in chart format. On the left side are displayed the interim criteria, those intended for application at program initiation. In the middle are the mid-term criteria, to be applied at the beginning of the second accreditation cycle. On the right side are listed accreditation criteria associated with the target accreditation program, the recommended mature form for the program. The accreditation criteria are divided into three groups, as a function of their purpose in the accreditation process. The three functions of criteria are to define mandatory features, to classify as to performance, and to rank according to LCC. The criteria associated with mandatory features are pass/fail criteria. The performance classification criterion will place an offered machine in its appropriate cell in the performance matrix, i.e., on its appropriate accreditation list. The LCC ranking criterion will determine the offered machine's rank within the appropriate accreditation list. The accreditation process would involve application of the criteria in the order described; mandatory features, performance classification, then ranking.

Interim	Mid Term	Target
	Mandatory Requirements	
Emulate ISAs of current standard computers	Emulate ISAs of current standard computers	ISA standardization at a common HOL level
Use current standard language	Use Ada HOL	Use Ada HOL
Validation of hardware compliance to emulated ISA standard	Validation of hardware compliance to emulated ISA standard	Validation of hardware compliance to Ada HOL/ISA standard
Require use of SEMs	Standardize at box level for accreditation	Standardize at box level for accreditation
Standard Interface between boxes	Standard Interface between boxes	Standard Interface between boxes
Two Thousand hour MTBF	Three Thousand hour MTBF	Four Thousand hour MTBF
Require use of SEMs	Require BIT to diagnose to module level	Require BIT to diagnose to module level
Require use of SEMs	Maintain and Spare at Module Level	Maintain and Spare at Module Level
	Performance Classification	
Use existing performance levels	Use accreditation performance matrix	Use accreditation performance matrix
	Ranking	
Consider LCC in evaluating candidates	LCC Model is major discriminator between candidates	LCC Model is major discriminator between candidates
	Accreditation Cycle Length	
Five years	Five years	Five Years
	Number of Computers per List	
Two in each performance range	Competition may limit number on list	Competition may limit number on list

5.2.1 Interim Criteria. The interim criteria are intended for use at the initiation of the accreditation acquisition policy. The recommended interim criteria, with only a few exceptions, follow the procurement scheme established for the Navy's current acquisition of the AN/UYK-44 successor to Sperry Univac's standard AN/UYK-20, and the AN/UYK-43 replacement for Univac's AN/UYK-7. In order to capture the Navy's large operational software base, some requirements must be levied regarding ISA and language for new computers. The desired result can be attained by specifying use of an existing ISA, which would restrict technological improvement, or by requiring the use of emulation to capture as necessary existing software. There will be a continuing need to validate candidate hardware to ensure that it operates as claimed with exising operational routines. A standard interface is vital, to assure operability within existing platforms and embedded systems. The current requirement for MTBF of 2000 hours is adequate. Requiring manufacturers to use Standard Electronic Module (SEM) boards in designing new computers in the interim will cause some restrictions to technology infusion, but it will simplify logistics at sea. Until more precise application categories can be defined, computers should continue to be classified into three performance levels as has been done in the recent past. Life-cycle cost will be one element in determining which candidate machines will be selected for use by the Navy. Currently available LCC modeling techniques are not sufficiently developed to enable LCC to be a major discriminator between candidate machines. The major difference between the initial accreditation program and the UYK-43/44 acquisition plan is a provision for accrediting the top two machines offered in each performance range. Having two competing machines in the inventory at the same time should stimulate the respective producers to keep their machines up-to-date. As Navy program managers seek computers for their embedded needs,

they will have a choice between two on-the-shelf offerings. The manufacturer which has the better over-all machine is likely to make more sales in such an environment.

5.2.2 Mid-Term Criteria. Accreditation criteria for the mid term period of the accreditation program will have undergone some progress reflecting changes in technology and development of evaluation tools. At this period in the evolution of the accreditation program, emulation of ISAs of current standard coputers will still be required to capture existing software. Use of the Ada language will be required for any new application software written for embedded applications. Validation of hardware will still be involved with checking compliance to emulated ISA standards. For purposes of accreditation, computers to be used in embedded applications must have a standard interface at the box level for operating with other application-system components. Progressing technology should support an increase in MTBF requirements to 3000 hours, thereby helping reduce logistics costs. Built-in test (BIT) capable of identifying faults to the module level will be required. Computers which are accredited at the box level will be maintained and spared at the module level. Information will be required from the manufacturer on fault populations in his product to serve as a guide to spares inventory level. A set of five application classes and three performance levels will have been defined, along with appropriate classification methodologies. The result will be 15 separate accreditation lists. A given embedded computer will likely meet the requirements for inclusion in more than one accreditation list. No upper allowable limit on the number of machines on a single accreditation list has been defined. However, competition between machine manufacturers may ultimately limit the number of machines. LCC modeling techniques should be

refined to the point where LCC will become the major discriminator between existing and new candidate machines on any specified accreditation list.

5.2.3 Target Criteria. The major change in criteria in moving to the mature accreditation program is the required standardization on a specified HOL/ISA. Validation will be involved with testing candidate machine compliance to the prescribed HOL/ISA standard. Based on a worst-case 90-day submarine patrol, a specified MTBF of 4000 hours should be adequate in consideration of the tradeoff between acquisition and logistic costs.

5.2.4 Summary. This set of accreditation criteria is not, nor is it intended to be a necessary and sufficient list of specifications required to conduct a computer acquisition program for embedded applications. For instance, the topic of using mil-spec, ruggedized or commercial standards was not addressed. The criteria examined and recommended are those peculiar to the concept of an accreditation strategy.

5.3 Transition Milestones

This section brings together the three sets of accreditation criteria with the factors affecting timing of program implementation, to produce a series of milestones for making the transition from the current policy of standardization to a mature form of an accreditation program. Figure 7 shows the resultant transition milestones. The plan indicates immediate implementation of the accrediation program, using the interim accreditation criteria. At the end of the first five year cycle, the mid term criteria will become effective for the next process of accreditation evaluation. By the end of the second five year cycle (1990), the necessary evaluation tools and advanced technology should support full implementation of the target accreditation criteria.

Figure 7

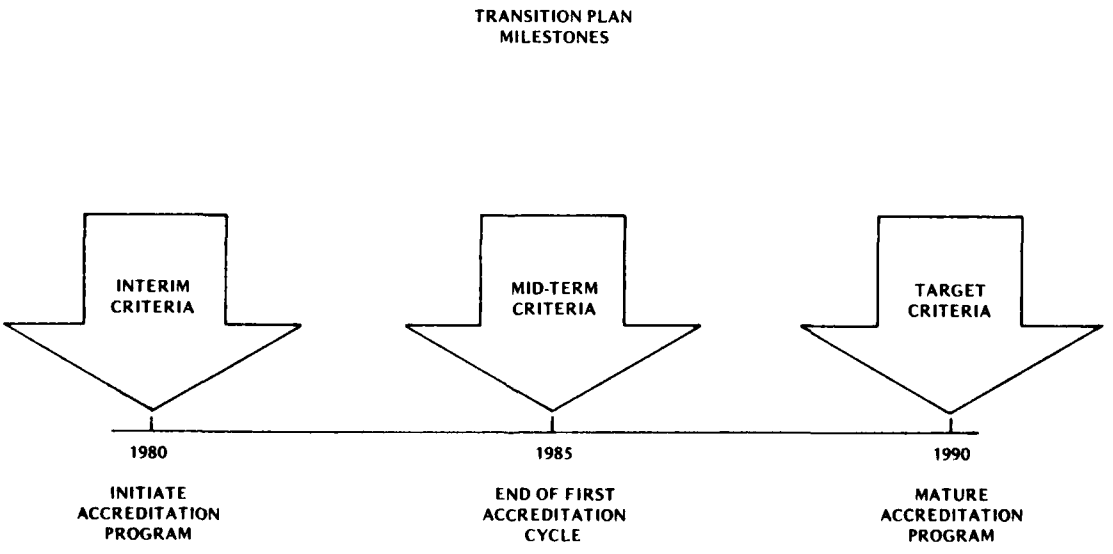

5.4 Government Leadership

The concept of accreditation as an acquisition policy is significantly different from the current Navy policy. Navy program managers and the computer industry will have to be educated on the concept of accreditation in order to gain their support. Key elements in gaining acceptance of a new concept are to have a well-defined policy and program, to announce to the world what that policy and program is, and to demonstrate sponsor support and interest in the program. Navy program managers have the objective of successfully bringing their system to Initial Operating Capability (IOC) and deployment. They must be shown how the accreditation concept will help them achieve their objectives. The computer industry is profit motivated and will produce equipment that will earn for them the greatest return on investment. However, the industry cannot react instantaneously to surprise demands and must be given guidelines as to what the demand will be well in advance of the requirement.

In a study of new generation computers by ITEK, it was reported that while there are definite trends towards commonality in new military computers, particularly in the physical and packaging aspects, there does not appear to be a real tendency towards reduction of computer proliferation and the continual development of new, unique designs. The commonality trends recognizable in these computers appears to be dead-ended at the "almost" common level of functionality. Thus commonality will not be carried to its logical conclusion, with the attendent cost benefits in support and maintenance, without strong DoD direction. The same trend is recognizable in the software aspect of these new generation computers. While computer organization and instruction sets are becoming similar, there is sufficient uniqueness between them to require completely separate software development

and maintenance tools for each machine. Furthermore, the trend towards microprogramming is generally not now producing computers that can be considered general-purpose emulators which can be used to provide the flexibility necessary to apply new common hardware technology to existing embedded system updates.[14]

Thus, the initial step in implementing an accreditation program is to announce the intention to do so, laying out the criteria and describing when they will become applicable. This step must be followed up with a continuing display of interest and support for the program, to ensure that the necessary evaluation tools are developed and that industry is responding with appropriate technology advancement.

6.0 SUMMARY AND RECOMMENDATIONS

This section briefly summarizes the effort and results addressed in the previous sections and presents recommendations regarding additional study needed to provide mature accreditation criteria and areas in which the Navy should exert its influence to promote acceptance and enhance the viability of the accreditation program.

6.1 Summary

The purpose of the study reported in this document was to examine the viability and strategy appropriate to a new, more nearly "optimal" acquisition policy called accreditation. Based upon the Navy's application needs for small general purpose computers and upon the requirements for competition and its resulting benefits, five application areas and three levels of performance have been defined, thereby establishing fifteen different accreditation lists. At least two machines per list are necessary to satisfy the goal of accreditation with regard to competition. Attempts to determine the upper limit in the number of machines in an accreditation list were unsuccessful because of insufficient information.

To move smoothly from current policy to accreditation, a time phased approach is necessary. To implement the policy, three sets of criteria have been derived based upon a combination of technological projection, parametric analysis of logistic costs and an industrial survey. It was determined that a five-year cycle is appropriate to accommodate a progressively more mature accreditation program. The initiation of the program should begin with the acquisition of the AN/UYK's 43 and 44 and reach maturity by 1990 with the target criteria requiring an HOL/ISA standard and a comprehensive life-cycle

cost model.

6.2 Recommendations

The iterim criteria listed in Section 5.2 can be implemented immediately. However, implementation of the mid-term and target criteria will require two types of additional effort. One type is the additional study of accreditation criteria and specifications. Included in this group are the following:

- o A detailed definition of application classes for Navy systems.
- o Development of benchmark routines or instruction mix equations for application classes.
- o Development of a comprehensive life-cycle cost model.
- o Determination of the upper limit on accredited machines.

The other type of effort required deals with Navy policy level emphasis on DoD programs and the implementation of Navy-wide efforts to accumulate and codify requisite data and to initiate programs to take advantage of the accreditation strategy. Included in this group are the following:

- o Support of DoD efforts in the Ada, VLSI and VHSIC programs.
- o Provide firm guidance to industry regarding requirement to develop and implement built-in-test and fault tolerant machines.
- o Development of a plan and initiation of a program to collect data in support of the life-cycle cost model elements.
- o Modify maintenance and training policy in accordance with the accreditation program.

Finally, the Navy must take an active leadership role. It must establish accreditation as its policy, educate its project managers as to its benefits, announce its plans and goals to the industrial community, and move to support the program through its fruition.

REFERENCES

1. "Final Report of the Navy Embedded Computer Review Panel," prepared for the Secretary of the Navy Office of the Assistant Secretary for Research, Engineering and Systems, October, 1978.

2. M. Fernandez, "Dod and NASA Computer Standardization Revisited," Computers in Aerospace Conference, 1979.

3. H. S. Stone. "Final Report Life Cycle Cost Analysis of Instruction-Set Architecture Standardization for Military Computer-Based Systems," prepared for U. S. Army Research Office, January, 1978.

4. E. M. Timmreck, "Computer Selection Methodology," ACM Computing Surveys, Volume 5, Number 4, December, 1973, pp. 199-222.

5. "New Generation Computer Evaluation for Military Computer Family (MCF) Implementation Plan," prepared for Army Electronics Command, February, 1977.

6. L. Bragg, H. Hylton, W. H. Kohler, H. S. Stone, "Life Cycle Cost Factors for MCF Logistic Scenarios," prepared for Army Research Office, December, 1978.

7. W. H. Kohler, H. S. Stone, "An Analysis of the Effect of Computer Standardization on Spares Inventory," prepared for Army Communications Research and Development Command, April, 1979.

8. Robert Shishko, "Discount Rate Selection for Defense Decision Making," Defense Management Journal, June, 1979, pp. 22-26.

9. J. B. Clary, R. A. Sacane, "Self-Testing Computers," Computer Magazine, October, 1979.

10. W. G. Bauricius, W. C. Carter, E. P. Hsich, D. C. Jessep, G. R. Putzolu, C. J. Tan, A. B. Wadia, "Cost Effectiveness of Self Checking Computer Design," IBM Thomas J. Watson Research Center, Yorktown Heights, New York, 1977.

11. L. Bragg, H. Hylton, W. H. Kohler, H. S. Stone, "Life-Cycle Cost Factors for MCF Logistic Scenarios," prepared for Army Research Office, December, 1978.

12. W. H. Kohler, H. S. Stone, "An Analysis of the Effect of Computer Standarization on Spares Inventory," prepared for Army Communications Research and Development Command, April, 1979.

13. Ibid

14. "New Generation Computer Evaluation for Military Computer Family (MCF) Implementation Plan," prepared for Army Electronics Command, February, 1977.

GLOSSARY OF TERMS

Accreditation	A strategy for the acquisition of general purpose mini or micro computers whereby a controlled number of computers meeting certain qualification criteria are approved by the Navy for use by project managers.
Accreditation Cycle	The period of time between opportunities for computer manufacturers to present their machines for accreditation.
ISA (Instruction Set Architecture)	The timing independent information about a computer that a programmer must know to write programs for that machine.
HOL (High Order Language)	A computer programming language that is machine independent.
HOL/ISA	The design of a machine's ISA such that a specified HOL may be executed directly and efficiently by that machine.
LCC	Life-Cycle Cost.
MCF	Military Computer Family (The Army's Embedded Computer Standardization Program).
NECS	Navy Embedded Computer System (A Navy [NAVMAT] program to develop a new computer for ship board embedded computer applications).
CMS-2	A Navy HOL.
J73, J731	An Air Force HOL, known as JOVIAL.
ATLAS	An IEEE HOL designed for testing purposes.
SEM	Standard Electronic Module, A Navy program to develop standard module for specified electronic function which are inherent to various Navy electro-mechanical systems.

Appendix A

FORWARDING LETTER, QUESTIONNAIRE AND ADDRESSES

Dear

It appears that the DOD must change its policy regarding the acquisition of embedded computers. The current policy of standardization on hardware has resulted in restriction of competition, inadequate injection of new technology and lengthy acquisition cycle time. It has been proposed that a policy of accreditation be adopted as a means of solving the problems associated with hardware standardization. In general terms, accreditation is a scheme whereby embedded computers are approved or accredited against a known set of criteria and placed on an accredited list, for use by DOD program managers in fulfilling their needs for embedded computers. Candidate computers would be evaluated periodically on the basis of performance, reliability, repairability and life cycle cost. Those machines meeting the established accreditation criteria would be added to the approved list and would remain on the list as long as their qualifications exceeded those of new candidates. A target criterion of the proposed accreditation concept is standardization upon a high-order language for all military embedded computer applications.

The Computer Science and Technology Laboratory of the Georgia Institute of Technology is examining the viability of some of the aspects of the accreditation concept. Through this letter and the attached questions, we request your and (several other companies') cooperation and participation in deriving some criteria which may be more generally acceptable to industry. Please take some time to answer the questions posed, and as you deem appropriate, add any additional thoughts pertaining to the subject.

An enclosed envelope is provided for the return of your answers. We would appreciate your response by February 29, 1980. Please be assured that your answers will be held in confidence and that no single response will be exposed. Again, the purpose is to obtain insight to consensus on accreditation criteria. Should you have any questions, please contact me by telephone at (404) 894-3464.

Yours truly,

H. B. Teates
Computer Science and
Technology Laboratory
Georgia Institute of Technology

HBT/jg

Enclosure

Instruction Set Architecture Standardization

An Instruction Set Architecture (ISA) is defined here to be all of the timing independent information about a computer necessary to write software for that machine. An ISA standard does not include instruction timing information or any implementation details not visible to the programmer (e.g. the existence of cache memory, number of memory parity bits, add time, multiply time, interrupt latency, etc.). It would include those details about privileged instructions, memory translation instructions, etc., necessary to implement operating systems and system software.

It is useful to decompose the structure of a computer system into a series of levels, ranging from the hardware circuit level to the ISA that the computer user sees. For the purpose of development of this concept, it will be assumed that the computer is microprogrammed with the microprocessor engine labeled as the level one machine. This level one ISA is used to implement a level two ISA through a microcode interpreter. This level two ISA is what is commonly referred to as the conventional machine language ISA. Most modern computers are implemented by emulation of the level two machine by a level one microcoded machine.

ISSUES

Studies in software engineering suggest that ISA standardization should be at the HOL machine level. Such a standardization policy would result in maximizing the robustness of software systems and allow the freedoms desired for technology infusion. Assume that the target accreditation factor would be ISA standardization at a common HOL level. Suppliers would be allowed to supply this ISA system at a variety of levels. They could supply a machine with an underlying ISA plus the necessary compilers, run-time systems, etc., to implement the HOL ISA, or implement as much of the HOL ISA at the level two machine as technology would permit. (Note: Standardization at the HOL level is impractical at the current time for many reasons, but can be recognized as a long range goal for interim planning).

Question 1: Do you feel that a HOL ISA standardization is a worthwhile and realizable goal?

Question 2: Do you agree that Ada is the HOL which should be adopted as the standard?

_____Yes _____No

Question 3: In what areas does Ada need to be extended before it could be adopted as the standard HOL?

It is recognized that there are certain requirements that must be met before an HOL ISA standardization effort is reasonable. Among these are not only the existence of a common HOL, but also freedom from dependence on traditional machine language programming. Also, for embedded computer applications, better methods for accessing underlying hardware (low level I/O) from HOLs are probably required before dependence on machine language can be eliminated. The latter will most likely require advances in software technology, even beyond the provisions attempted in the Ada effort. Nevertheless, HOL standardization will require continuing or increased pressure on program managers to use HOLs in project development. That the DoD is moving toward common HOLs is evidenced by directive 5000.31 and the current Ada language effort.

Question 4: Assuming that HOL standardization at the ISA level is desireable, do you feel that the movement toward the use of HOLs is progressing with adequate speed? Are the current directives adequate? Should they be strengthened? Are they too restrictive?

Question 5: What additional technical problems need to be solved before an HOL ISA standardization policy is reasonable?

Question 6: How long (in years from now) do you estimate the period to be before an HOL embedded computer standard would be appropriate?

Any move away from the current military policy of standardization must not compromise the ability of the military to fulfill its mission. It is for this reason, as well as minimization of life cycle costs, that an accreditation policy must be established that captures existing software. Operational military software systems depend heavily on the machine language architectures of the current standard computers. It is estimated that over half of the existing software is written in machine language for the UYK-7, GYK-12, AYK-14, UYK-19, UYK-20, etc., architectures. Therefore, the accreditation criterian must include, in the near term, standardization on these instruction set architectures. It is a worthwhile goal to aim for an HOL standardization as an ultimate objective, however, until the software base of the military can

be captured by that HOL commonality, it remains a future target.

It is important that requirements not be placed on detailed instruction timings. In order to ease technology insertion, machine designers should have the freedom to trade off architectual elements. The subject of the definition of performance ranges will be addressed in the next section. Part of the accreditation process must necessarily be a set of ISA verification programs and procedures to validate the compliance of a particular piece of hardware to the ISA standards. These programs should be quite similar to the diagnostics commonly provided by computer manufacturers to determine and isolate hardware problems.

Question 7: Do you feel that this is a reasonable requirement for accreditation? Is the relaxation of detailed timing specifications sufficient to allow technology insertion to the degree manufacturers desire?

The ISA requirement should be an absolute one with subsetting and supersetting forbidden. The reason for no subsetting is to capture existing software bases. Supersetting should be disallowed to prevent non-controlled proliferation of similar but distinct architectures. It is certain that if enhancements to an ISA exist, use of those enhancements will exist, thus nullifying the transportability across members of accredited systems. This is not to say that as the standardization evolves toward an HOL ISA, new ISA

standards which superset old ones will not exist; but rather, that periodically the ISA standard should be updated in a controlled way (presumably supersetting the old standard) and this new standard will be strictly enforced in the accreditation program.

Question 8: Do you agree that subsetting is undesireable? How about supersetting? How often should ISAs be reviewed to allow advancements in technology? Explain.

Question 9: What do you believe to be (in years) the time period that should be allowed between accreditation cycles to allow a significant change or advancements in technology in ISA improvement?

5 years 8 years 10 years 15 years

Performance Factors

BACKGROUND

The criterion of a standard ISA for accreditation avoids the subject of performance or instruction timing by design. Procedures and criteria must be established to classify computers into performance ranges for addition to an accredited list. This list of accredited computers will actually consist of a number of lists, one for each performance range. All members of all lists will conform to a fixed ISA definition. We are suggesting that the categorization of machines into performance levels be accomplished by establishing benchmark techniques that measure instruction mixes common in actual field uses. Assuming that such procedures can be established, there will be a relaxation of the requirements that specific instructions meet fixed timing thresholds. Therefore, computer designers will have some flexibility in trading off the speed of some instructions for others, maintaining an instruction mix bandwidth. Obviously, field applications which make use of ISA peculiarities (e.g. undefined instructions) or current ISA timing characteristics will not be guaranteed to be portable. However, programs which use such "quirks" will be considered to be in error even though they may execute on current machines.

[illegible] an interim ISA standard requirement for accreditation at the [illegible] the language level, performance measurement can be established

in the form of equations involving specific instruction mixes rather than the more traditional method of establishing benchmark programs. Such specifications should aid the computer designer in producing a machine which would execute in the performance range that it was targeted for. It is likely that the accreditation list will actually consist of multiple parallel lists, each characterized by an instruction mix performance equation. Each of these sublists will represent a different category of application. It will be expected that accredited computers will have entries in each of these parallel lists, depending on the established performance range of that machine. The determination of performance categories in each of the lists will be guided by military requirements. There will likely be a minimum of two entries in each list as a result of requirements for competition.

Question 10: Does the above definition of performance allow sufficient freedom in designing ISA emulators to be worthwhile? It is assumed that the above definition of performance will allow capturing the existing software base. Do you see any reason why this would not be the case?

Question 11: It is suggested that accreditation criteria will address four areas: performance, reliability, repairability and life cycle costs. Would you add any additional areas? What?

Question 12: Is throughput a sufficient measure of performance (i.e. in thousands of operations per second (KOPS)? What would you add?

Question 13: Would you divide the range of performance into categories such as:

100-300KOPS, 300-600KOPS, 600-1000KOPS?

Level of Standardization

One of the areas of controversy when considering a standardization strategy involves the level at which standardization should occur. The alternatives discussed here are card, module, or box level. Most LCC models indicate that sparing should occur at the card or module level, but that is not an issue here. What is being considered is the acquisition strategy leading to accreditation criterion. With a box level standardization, functionality of a complete computer would be specified. The number and kinds of subunits contained in that box would only be germane to the LCC evaluation of the box from a logistics, sparing, maintainance, training, etc., point of view. Module level standardization would require that modules acquired from different sources be plug compatible in some sense, with interchangability within a single box. Standardization at the module level seems attractive at first glance since procurement would take place at potentially the same level as sparing. Competition could take place at this level and the logistics problem is greatly simplified. Examples of attempts at standardization at the module level include NECS, SEM, and the initial MCF concept.

In order to standardize at the module level, form/fit/function constraints are required at the module level. In addition, in order that the modules be plug compatible, a standard bus definition is required. Proponents of module level standardization point to the existence in the commercial marketplace of

second source suppliers of memory and peripherals for existing machines. Given the existence of emulation as today's technological answer to upward compatibility and ISA implementation, module standardization is quickly extended to require compatibility across different ISAs. Examples are the NECS and MCF programs.

Proponents of box level standardization are quick to point out that module compatibility across multiple ISAs has yet to be proven. The argument is that to achieve efficiencies required by existing ISAs, the designer must not be forced into conforming to a fixed bus standard, or for that matter to any other forced partitioning of the components of the system. The assertion is that module level standardization will stifle any large advances obtainable through technology improvement as well as remove incentives for industry to participate in such developments. If one looks at computer architecture implementations in industry, one tends to find large variations in bus architectures, even between different performance members of the same ISA family, e.g., the PDP-11 family.

Question 14: At what level should standardization occur: module or box? Is it possible to standardize on a bus and maintain enough design freedom to incorporate technology advances in new implementations of a standard ISA?

Once a newly accredited number is qualified, it might be possible to recompete the modules that make up that member, on a form/fit/function basis. This could not be a "build to print" type competition unless the government

owned the design. One of the consequences of such a policy is that profit incentives are reduced or removed for industry to participate in the R&D required to design new accredited computers.

Question 15: Given that competition will occur for inclusion in the accredited lists, is it desirable to separately compete modules that make up the box?

Question 16: Are you aware of any periodicity in the introduction of major changes in computer architecture? If so, about how long is the period?

Question 17: How soon do you expect the general implementation of major changes in computer architecture?

eg.:

---- language directed architectures

---- self-defining data

---- lexical-level addressing

---- variable-size storage cells

others:

Question 18: What factors, if any, currently are pressing for the introduction of new computer architectures?

Questions Regarding Maintenance of Embedded Computers

Question 19: What factors in computer design would facilitate the repair of computers in combat situations, in environments where spare parts are difficult to obtain beyond a limited, local supply, or where the repair technicians have minimal qualifications?

Question 20: If you subscribe to the philosophy of "design for repair", what are the primary methods you use to achieve such design?

Question 21: Comment on providing information to the government concerning:

fault populations -

mean time between failure -

mean time to repair -

skill level to repair -

Question 22: At what levels do you currently use Built-In-Test (BIT)?

---- chassis

---- module

---- board

---- IC

Question 23: What forms of BIT are you currently using or considering using?

---- error detecting and correction codes

---- hardware redundancy (replication)

---- resident software

---- firmware:

---- test patterns

---- test reference data

---- microdiagnostics

---- other (please specify):

Question 24: Do you plan to increase your use of BIT?

Question 25: What might be some major obstructions or objections to BIT?

Question 26: What problems might arise in implementing BIT in a multi-vendor, accredited system?

Question 27: What problems will need to be avoided or overcome in the communication of fault information in modularized computers, considering the need to preserve the identity of the fault source, the type of error occurring, and perhaps the state of the machine?

Questions Regarding Cost

It has been suggested that one of the accreditation criteria be cost. For example, the maximum number of vendors with accredited machines in a performance range may depend upon the amount of money available to support the logistics and repair of the machines. Using all contractor supported repair, costs would include both fixed costs and variable costs to perform repair. Fixed costs include test equipment, training, documentation necessary to perform logistical repair. Variable costs relates to the costs incurred to perform a specific module/box repair. To obtain a feel for thse costs, we would like to gather fixed and variable cost data for a computer like the AN/UYK-20. Therefore, the following questions:

Question 28: What are your estimates of the fixed costs?

Question 29: What are your estimates of the variable costs?

Question 30: What percent error would you allow for each of your answers to the above?

Another aspect associated with the concept of accreditation is how to determine if a new machine should be added to the accreditation list? Assuming that cost is one of the criteria to be used in dealing with this problem, it is necessary to have a feel for the following type of information:

(1) Investment dollars (R&D) required to field a new machine.

(2) Reductions in cost from using the new machine. Reductions may be in terms of volume, weight, power, maintenance, spare parts, purchase price, software costs and so on.

To obtain a feel for these costs, we would like to gather data for a computer like the AN/AYK-14. Therefore, the following question?

Question 31: What are your estimates of the investment required as a function of time to field a new machine?

Question 32: Do you believe current life cycle cost models are adequate for the job envisioned? If inadequate, what are the cirtical aspects that must be developed?

<addressee>Emmett W. Johnson
<address>Sperry Univac Defense Systems
Univac Park, P. O. Box 3525
St. Paul, MN 55165
<>

<addressee>Walter Loewenstern
<address>Rolm Corporation
8027 Leesburg Pike, Suite 107
Vienna, VA 22180
<>

<addressee>Leon Bloom
<address>Litton Data Systems
8000 Woodley Avenue
Van Nuys, CA 91409
<>

<addressee>
<address>Harold L. Ergott
Noden Systems
Norwalk, CT 06856
<>

<addressee>W. L. Wolles
<address>Control Data Corporation
3101 East 80th St
Mailing Address/Box 609
Minneapolis, MN 55440
<>

<addressee>Joel J. Buschek
<address>Booz Allen Applied Research
4330 East West Highway
Bethesda, MD 20014
<>

<addressee>Dan Kober
<address>Arinc Research Corporation
2551 Riva Rd
Annapolis, MD 21401
<>

<addressee>Bert Courtang
<address>Teledyne Systems company
19601 Nordhoff St
Northridge, CA 91326
<>

<addressee>J. M. Bradley
<address>Data Processing Product Division
Hughes Aircraft Co
Fullerton, CA 92634

<>

<addressee>G. B. Stearns
<address>Avionics Division
Honeywell, Inc.
13350 U.S. Highway 19
St. Petersburg, FL 33733
<>

<address>The S-1 Project
Lawrence Livermore Laboratory
University of California
Livermore, CA 94550
<>

<addressee>B. Casey
<address>Lear Siegler, Inc.
4141 Eastern Ave S.E.
Grand Rapids, Michigan 49508
<>

<addressee>Henry Nye
<address>Digital Equipment Corp.
200 Forest St.
Marlboro, MASS 01752
<>

<addressee>Casmir Barcynski
<address>Burroughs Corp.
Special Systems Division
Box 517
Paoli, PA 19301
<>

<addressee>Herman Fisher
<address>Litten Data Systems
8000 Woodley Ave
Mail Stop 64-48
Van Nuys, CA 91409
<>

<addressee>H. C. Turner, General Manager
<address>Computer and Instrumentation Division
Westinghouse Electric Corp.
1200 West Colonial Dr
Orlando, FL 32804
<>

<addressee>Edward Spignese
<address>Ratheyon Equipment Division
600 Spring St.
North Dighton, MASS 002764
<>
<addressee>Tom Sleight

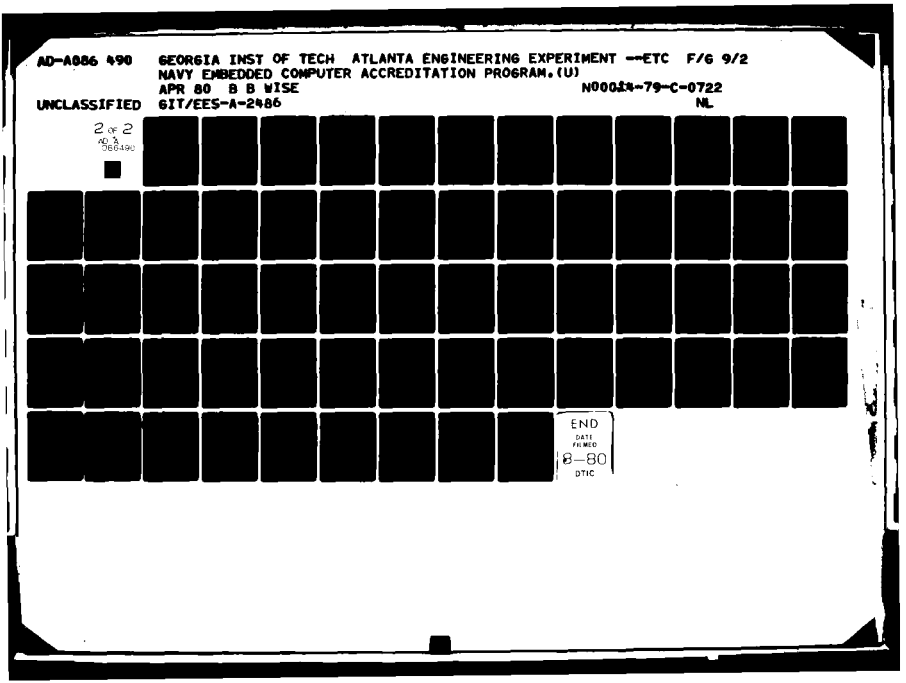
AD-A086 490
GEORGIA INST OF TECH ATLANTA ENGINEERING EXPERIMENT --ETC F/G 9/2
NAVY EMBEDDED COMPUTER ACCREDITATION PROGRAM.(U)
APR 80 B B WISE
N00014-79-C-0722
UNCLASSIFIED
GIT/EES-A-2486
NL
2 OF 2
AD A086490
END
DATE FILMED
8-80
DTIC

<address>Applied Physics Laboratory
John Hopkins University
John Hopkins Rd
Laurel, MD 20810
</>

<>

<addressee>G. B. Stearns
<address>Avionics Division
Honeywell, Inc.
13350 U.S. Highway 19
St. Petersburg, FL 33733
<>

<address>The S-1 Project
Lawrence Livermore Laboratory
University of California
Livermore, CA 94550
<>

<addressee>B. Casey
<address>Lear Siegler, Inc.
4141 Eastern Ave S.E.
Grand Rapids, Michigan 49508
<>

<addressee>Henry Nye
<address>Digital Equipment Corp.
200 Forest St.
Marlboro, MASS 01752
<>

<addressee>Casmir Barcynski
<address>Burroughs Corp.
Special Systems Division
Box 517
Paoli, PA 19301
<>

<addressee>Herman Fisher
<address>Litten Data Systems
8000 Woodley Ave
Mail Stop 64-48
Van Nuys, CA 91409
<>

<addressee>H. C. Turner, General Manager
<address>Computer and Instrumentation Division
Westinghouse Electric Corp.
1200 West Colonial Dr
Orlando, FL 32804
<>

<addressee>Edward Spignese
<address>Ratheyon Equipment Division
600 Spring St.
North Dighton, MASS 002764
<>
<addressee>Tom Sleight

<address>Applied Physics Laboratory
John Hopkins University
John Hopkins Rd
Laurel, MD 20810
<>

Appendix B

SUMMARY OF REPLIES TO NAVAL ACCREDITATION STUDY QUESTIONNAIRE

SUMMARY OF REPLIES TO NAVAL
ACCREDITATION STUDY QUESTIONAIRE

Question 1: Do you feel that a HOL ISA standardization is a worthwhile and realizable goal?

5-yes
0-no
1-split

One of those answering "yes" qualified his answer with the provisos that the standard must handle target dependencies in all HOLS, must permit code insertion, and must be free to evolve.

The split answer asserted that such a goal is indeed worthwile, but that the "software environmental tools should not be developed and standardized until the language is clearly defined." Whether or not the goal is realizable would depend on whether or not the language and support software can be self-hosted on the target machines (e.g., micros) without jeopardizing mission success. This respondent seemed to feel that self-hosting was necessary.

Question 2: Do you agree that Ada is the HOL which should be adopted as the standard?

4-yes
1-no
1-maybe

The respondent answering "no" indicated that it is too early for Ada, but that Ada would be appropriate when matured and available. The one answering "maybe" wanted to wait for an Ada compiler to see what impact the language will have on timing and memory constraints.

Although the response was quite favorable toward Ada, half of those responding also wanted other accredited HOLs, such as F77, FIV, CMS-II, J73, J73I, and Atlas. One of these said that any language meeting DoD I-5000.31 should be available, as well as future viable languages. There seemed to be some confusion or disagreement over having a single HOL standard as opposed to having a list of accredited HOLs.

Question 3: In what areas does Ada need to be extended before it could be adopted as the standard HOL?

One respondent said an application study would be needed to determine the answer. Another indicated that there should be no extensions, although minor modifications might be acceptable. Two thought that subsets should be allowed. Other suggestions are listed below.

HOL constructs for communicating with a broadcast bus and other shared communication paths are needed, due to expectations of multiple processor systems with up to hundreds of processors.

Need access to underlying hardware.

Revised syntax for embedded assembly code to follow syntax and operations format of underlying ISA.

Ability to pass procedures as parameters.

Low-level system programming capability.

Low-level I/O programming capability.

(Demonstrated) efficient multitasking, context.

Switching and optimized code (with regard to both time and space).

In addition to extensions to the language, it was suggested that a HOL standard requires change in hardware and firmware design to remove I/O details from the software. An example given was Litton's GYK-12 emulator for interfacing peripheral I/O, which transparently reblocks messages.

Question 4: Assuming that HOL standardization at the ISA level is desirable, a) do you feel that the movement toward the use of HOLs is progressing with adequate speed? b) Are the current directives adequate? c) Should they be strengthened? d) Are they too restrictive?

a) 2-yes
3-no
1-indecipherable

Of those answering "yes," there was concern that directives might be made prematurely (e.g., FORTRAN66 vs FORTRAN77). Software environmental tools should be developed and tested first.

Of those responding "no," one wanted the use of directives to speed up the movement, while another was against the use of directives, saying they wouldn't work and preferred to see a good product (e.g., Ada, compiler) which would create a desire on the part of project directors to use the HOLs. A third respondent expressed concern about the Navy's "Who, me?" attitude toward support of Ada.

b) 1-yes
3-no
2-indeterminable

One felt that the current directives are premature.

c) 3-yes
2-no
1-indeterminable

d) 3-yes
1-no
1-no comment

One of those responding "yes" felt that there was a need for more flexibility in the application of existing directives.

Although there was no general agreement as to what should be done about directives, everyone seemed discontent with them.

Question 5: What additional technical problems need to be solved before an HOL/ISA standardization policy is reasonable?

None

Acceptance: "Standards follow, not lead, user acceptance."

Transportability to host computers

Metrics

Testing techniques

Application to process of designing application programs

Reasonable execution speed

Machine architecture complimentary to HOL

Compiler availability

Compilers must be in public domain (unlike CMS-II)

Definition of minimum hardware required to support HOL standards

Definition of allowable host machines

Dealing with the need (if any) to self-host (When will microprocessors and associative memory systems permit practical self-hosting?)

Access to software tools on tri-service basis

Question 6: How long (in years from now) do you estimate the period to be before an HOL embedded computer standard would be appropriate?

3-now
1-now or in 1981
1-in 1985
1-in 1985-86

One respondent said that such a standard would be appropriate now, but that it would not be enforceable until Ada compilers (with optimizing) became public domain. Another indicated a need now for a better design notation from which an appropriate HOL standard might evolve.

Question 7: a) Do you feel that this is a reasonable requirement for accreditation? b) Is the relaxation of detailed timing specifications sufficient to allow technology insertion to the degree manufacturers desire?

a) 5-yes
1-no

One "yes" answer was qualified by a requirement that it be properly managed for competition.

b) 1-yes
2-no
3-other

One of those answering "no" said that there is also a need for "relaxations on modularity, operator interface improvements (human engineering), and inclusion in the ISA's a capability to effectively control fault tolerant features."

Although one argued for no timing specifications at all, two respondents said that specifications for minimum timing should be set while allowing faster times. Another felt that there is a need to develop new hardware for direct execution of the HOL and that "unless I/O is emulated identically even though deviced upgrade (sic) by state of the art, there is little that can be captured of whole programs."

Question 8: a) Do you agree that subsetting is undesirable? b) How about supersetting? c) How often should ISA's be reviewed to allow advancements in technology? Explain.

a) 1-yes
4-no
1-wait and see

One respondent said that subsetting should be allowed but must be controlled. Others said subsetting is needed for special cases and is necessary if Ada is forced on existing target machines.

b) 1-yes
3-no
1-indeterminable answer
1-should be optional

The respondent agreeing said that supersetting is highly undesirable and counter to the basic concept of standardization. A dissenter felt that it is mandatory for new microcoded functions, etc., which arise due to new applications. (Consider whether this might not be handled in Ada through the use of modules.)

c) 6 mos. - 1 year
1 - 2 years
2 years
5 years

Question 9: What do you believe to be (in years) the time period that should be allowed between accreditation cycles to allow a significant change or advancements in technology in ISA improvement?

1 year
3 years max
5 years (from two respondents)
8 years

None - accreditation cycles are driven by technology improvements; review annually or biennially

Question 10: a) Does the above definition of performance allow sufficient freedom in designing ISA emulators to be worthwhile? b) It is assumed that the above definition of performance will allow capturing the existing software base. Do you see any reason why this should not be the case?

a) 3-yes
1-no
2-no answer

Comments suggested that the mix must include interrupt response timings, task/context switching, I/O setup and performance, and memory management in setup, context switching and access; otherwise, a benchmark would be needed. Also, the mix speed validation test program should be in the public domain.

b) 3-yes
2-no
1-perhaps

It was questioned whether the existing code is really worth capturing, and it was pointed out that one can't assume that the existing software base is error free, which would be necessary for certain capture. An alternative suggested was to translate from ISA to ISA rather than to emulate. It was noted that emulators are hard to check out and verify, especially for exception conditions. There was concern that it might not be possible to capture time dependent software; two respondents argued that each instruction must run as fast as the standard and that the time and sizing for each instruction must be measured.

Question 11: It is suggested that accreditation criteria will address four areas: performance, reliability, repairability, and life cycle costs. Would you add any additional area? What?

Existence of a commercial equivalent

Peripherals

I/O

Support software

Interface accreditation (particularly to data base)

Each instruction and number calculation are in accordance with ISA standard (use tests here)

Maturity of candidate

Development cost

ISA conformance to capturability criteria

Size

Weight

Power

Security

EMI

Radiation hardening

Expandability

Fault tolerance

Delivery schedule and quantities

Unit cost

MIL spec level

Financial guarantees (escrows) of life cycle cost projections

Availability

Survivability

Repairability "quantified" to level needed (card, box, chip,...)

Second-sourceability

Life cycle guaranteed availability of chips, components, mechanical items, testers, and support software host equipment

Provision to measure repairability at card level, especially where manufacturer proposes to do this and especially if card replacement is repair concept in performance equations

Volume

New ones to be better than prior generation computers, since such improvements might open up new applications.

Question 12: a) Is throughput a sufficient measure of performance (i.e., in thousands of operations per second (KOPS)? b) What would you add?

a) 2-yes
4-no

b) I/O bandwidth
I/O benchmark

Assurance of instruction and number calculation performance

Task-switching

I/O setup and performance

Memory management setup and performance

Diagnostics thoroughness

Weight, volume, and power consumption as a function of throughput and memory size required

For I/O: single channel minimum rates (per channel type); aggregate minimum rates; guaranteed minimum rates per channel type per channel priority

Interrupt latency and channel switch times (from event to execution of first instruction after state switch)

KOPS must be based on a given mix and given operands

Data storage bandwidth hierarchy

Question 13: Would you divide the range of performance into categories such as: 100-300KOPS, 300-600KOPS, 600-1000KOPS

1-yes
5-no

The one answering "yes" volunteered the following scale:
10-100KOPS
100-300 KOPS

300-600KOPS
600-1000KOPS
1000-2000KOPS
2000-4000KOPS

Alternatives were offered by those answering "no":

Rather have a "continuous" scale

Measure mixes on existing generation computers (e.g., 370, VAX, UYK-20, AYK-14, UYK-14, GYK-12) to see what execution rates really exist; then for each ISA, specify a full speed and a half speed machine

KOPS numbers need to be based on empirically measured performance on real present or prior generation hardware

Experience shows that with KOPS requirements and non-public domain GFI timing programs, pressures of competition lead to operationally unrealistic "tuning" of measurement programs

Mission success depends on more than throughput - weight, volume, power consumption as a function of throughput, memory size needed, reliability, and life cycle costs: the program manager's viewpoint as opposed to the programmer's viewpoint

Question 14: a) At what level should standardization occur: module or box? b) Is it possible to standardize on a box and maintain enough design freedom to incorporate technology advances in new implementations of a standard ISA?

a) 0-module
5-box
1-indeterminable

It was commented that standardizing at the box level simplifies configuration management and control and would not require "computer integration" contracts. Also a single supplier would be responsible for operations and support of the complete computer system. It was further suggested that the box should be defined functionally and by interconnects, but not by size. "For box standards like ISA standards, each service would be responsible for box dimensions, but the standard would be responsible for functionality and interconnects." One of those choosing the box level of

standardization said that the module level such as memory and power supplies has proven acceptable.

b) 0-yes
4-box
2-no answer

An alternative suggested was a family of bus architectures. One respondent said that internal bus saturation often is the limit on throughput. One of the two not answering the question directly said, "Bus standardization is technology sensitive and should not be frozen over a longer period than the lifetime of a particular accredited box."

One response to the entire question was "Define module interfaces; define box interfaces. Then competition on an announced 5-year cycle to capture new technology. Logistics simplified since interfaces are held constant."

Question 15: Given that competition will occur for inclusion in the accredited lists, is it desirable to separately compete modules that make up the box?

1-yes
4-no
1-alternative suggested

The one answering "yes" suggested letting the vendors keep the design as proprietary, but forcing them to conform to interfaces and recompeting on a regular cycle. One problem pointed out is that if module level standardization is chosen, who decides which vendor's module is at fault when a combination doesn't work, even though all meet specs? One "no" answer allowed for the exception of large modules, e.g., SEMS-9 memories, which have very simple interfaces, high dollar content, and existing multiple sources. Another was against separately competing modules unless full development and tooling is paid to multiple suppliers.

The alternative suggested was to pursue accreditation of a design first, then second sourcing of modules by either private competition (with the contractor retaining the card design proprietarily) or public competition (if the government owns the design and rights). "Module competition should be a result of and not a preexisting condition before accreditation."

Question 16: a) Are you aware of any periodicity in the introduction of major changes in computer architecture? b) If so, about how long is the period?

a) 3-yes

3-no

Due to the sudden change in terminology (from "ISA" to "computer architecture"), there seems to have been confusion as to what was being asked for in this question. Most seem to have taken this to include hardware level architectural changes too. One respondent, who apparently recognized this as a question about instruction set architectures, said that they are the same now as they were fifteen years ago. It was pointed out that there is generally a need to allow "graceful evolution" of ISAs and associated software to protect "customer bases of business."

b) 2 - 3 years currently

3 years "due to LSI and logic type availability"

2 - 4 years "due to processor chip advances, I/O architecture and peripherals evolution, memory and memory chip advances, power supply and power source evolution, and unpredictable bright ideas

3 - 4 year randomized intervals for major advances

Question 17: How soon do you expect the general implementation of major changes in computer architecture?

1 - 5 years

e.g.:

Language directed architectures

now (2 respondents)

more than 5 years

6 years

about 10 years

in the 1990's

Self-defining data

never

2 years

more than 5 years

in the 1990's

Lexical level addressing

4 years

more than 5 years

in the 1990's

Variable size storage cells

now

more than 5 years

6 years

in the 1990's

15 years

Others:

Large associative memories - 10 years

Broadcast bus/receiver selection - 5 years

Greater instrumentation coverage - 2 years

Multi-address machines - 2 years

I/O architecture - 7 years

Interrupt and context switching - 10 years

Operating system design - 7 years

Memory architecture and addressing - 5 years

Interactive or transaction driven control technology - 3 years

The respondent consistently answering "in the 1990's" said that these are feasible now, but that the timing of the general implementation is a compatibility-based issue.

Question 18: What factors, if any, currently are pressing for the introduction of new computer architectures?

Data flow

Holographic memory

Array processors

Associative memories

Systolic arrays

Software costs

HOL compatibility

Common coherent interconnect scheme for multiple processors

Greater application coverage, e.g., signal processing functions and image processing functions

Hardware cost, performance, and reliability have improved dramatically in the last 20 years and give a basis for new tradeoffs.

Security/protection mechanisms

Adaptability to special processing functions and customized configurations

Flexible I/O

Size, power, and weight restrictions

Increased processing speed

Complexity of threat

"Crummy compilers"

NIH

"Better commercial architectures causing greener-pastures effect and designer's lust"

Advancing chip technology

Advancing memory technology and increased memory capacity

Peripherals and peripherals control architecture

Increasing emphasis on multi-user, transaction-driven applications

Military applications

Question 19: What factors in computer design would facilitate the repair of computers a) in combat situations, b) in environments where spare parts are difficult to obtain beyond a limited, local supply, or c) where the repair technicians have minimal qualifications?

a) Fail-safe capability

Multiprocessors

Automatic reconfiguration as result of built-in-test

Specifying fault tolerant hardware

Self-repair where practical

b) Standardized interfaces to permit use of old, new, and cannibalized parts

Spare at largest module level

Investment in component, box, and system design RMA

c) Repair at depot

Inclusion of an intelligent maintenance processor to aid in fault detection and isolation

Provide for spares plugged into basic unit

Adequate documentation

Super-good diagnostics, including at the card level, which exercise chips, circuit paths and cards instead of exercising instruction repertoire at functional level

Plug in chips

Extended built-in-test

Functional partitioning

Standard board sizes

Designing for two level rather than three level maintenance

Question 20: If you subscribe to the philosophy of "design for repair," what are the primary methods you use to achieve such design?

Self-diagnosis

Easy access to all replaceable units

"Modern software with its emphasis on software integrity and more readily maintainable (HOLs) will reduce repairs and lower costs of repairs"

Continuous asynchronous built-in-test

Standardized interfaces

Built-in-test to replaceable element, tests at depot level to point to bad component

Requirements and design reviews

Functional organization of machine partitioned by cards: Software and firmware fault diagnosis can isolate functional paths and logic functions which lead to recognizable replacement modules; then microdiagnostics, using firmware, access individual gates, paths and circuits in testing

Numerous test points at box level

Reset lines for sequential logic

Tap-in feedback loops to connectors

Tap-in free running clocks

Question 21: Comment on providing information to the government concerning: a) fault populations, b) mean time between failure, c) mean time to repair, and d) skill level to repair.

a) High temperature stressed units due to lack of cooling capability

Currently provided (two respondents)

Automated fault insertion tester, which shorts and opens every circuit exercising every diagnostic on each inserted fault and categorizing the results

The government must establish the fault class and ground rules

b) Estimates currently provided (two respondents)

Calculated numbers don't always match failure rate

Temperature is the limiting factor to producing 5k to 10k hour MTBF systems

Meaningless, as numbers are "engineered" differently by each major vendor...MTBF = chip designers * MIL handbook numbers

Current method OK

c) Estimates currently provided (two respondents)

Hard to measure on a prototype

Dependent on logistic chain

20 minutes

"A real world empirical number provable in tests"

Could be more meaningful if MTTR calculations included all repairs, not just those limited to ORG LEVEL repairs

Good thing to keep in acceptance tests

Current objectives unrealistic

c) Box level sparing combined with good built-in-testing and test procedure standards required as part of accreditation program will result in low skill level needed for repair

Should be low skill level for most failures; should not need to repair often -- why can't ship processors be as good as deep space processors?

9 - 3 = high school level

Currently provided (two respondents)

Need better definition of skill levels in MIL standards

Need a consistent definition among services

Question 22: At what levels do you currently use built-in-test (BIT)?

Chassis

5-yes
0-no
1-no answer

Module

5-yes
0-no
1-no answer

Board

2-yes
3-no
1-no answer

IC

1-yes
4-no
1-no answer

Question 23: What forms of BIT are you currently using or considering using?

6 - error detection and correction codes

4 - hardware redundancy (replication)

5 - resident software

Firmware:

5 - test patterns

5 - test reference data microdiagnostics

Other:

Hardware reconfiguration

Separate but built-in maintenance processor

Card self-tests designed into each card's logic

Off-line test multiplexers

Dead-man's timer

Wrap around signals

Question 24: Do you plan to increase your use of BIT?

5-yes

0-no

1-no answer

Comments:

Yes, but only as the specifications require and the development, production, and life cycle costs allow.

Yes, as much as possible. It's a real competitive edge. That's how we won (...contract name deleted here...) (among other factors).

Question 25: What might be some major obstructions or objections to BIT?

Cost - but not significant

Increased failure rate - but not significant

Front-end expense means losing the competition

Uses additional hardware, software, memory, and a small percentage of throughput

Cost - development and production

Added failure nodes of the BIT structure can degrade overall reliability to a minor degree

Expensive - customer funding difficult often

Diagnostics and microdiagnostics often cost much

Question 26: What problems might arise in implementing BIT in a multi-vendor, accredited system?

BIT procedures must be uniform for vendor independence to allow for low-skill technicians

Compatability - need well defined, standardized interfaces

Adequate procurement specifications for defining BIT capability

"Each vendor's BIT is different. Competition forces some to lie unless government insertion of faults is part of test."

"The responsibilities of each vendor must be defined and a method for revision must be established prior to issuing any hardware contracts."

Other considerations:

- Design responsibility
- "Build-to-print" level of detail necessary
- Non-standard interconnect structures
- Timing
- Module interface documentation
- Coordination of ECP's
- System design

Question 27: What problems will need to be avoided or overcome in the communication of fault information in modularized computers, considering the need to preserve the identity of the fault source, the type of error occurring, and perhaps the state of the machine?

Minimal problem if isolate to box level or module level if system has small number of modules

Need a coherent interconnect scheme

Fault message produced by BIT of failed module or by cooperating adjacent module; failure monitor accepts message and takes appropriate action

Reconfiguration is desirable to enhance survivability and availability

Coherent interconnect and redundant processors needed; solve the general case, not hundreds of specific software loads, etc.

Type of error occurring and state of machine needed for system development rather than in operation. (But how can left-over bugs be caught and improvements made during deployment otherwise?)

Faulty module obscuring correct fault module identity

Cost allowances for development and production

"We've been successfully doing it for years. IBM's been doing it since 370 days. Why problems with 10-year-established technology?"

Diagnosis of I/O errors requires additional hardware to the extent that it may not be cost effective for the more complex I/O channels. This problem should be studied in detail and hopefully a policy towards diagnosing I/O failures can be established.

Questions 28 - 30: a) What are your estimates of the fixed costs? b) What are your estimates of the variable costs? c) What percent error would you allow for each of your answers to the above?

(These are summarized here per respondent.)

First respondent:

"Cannot give general answer."

Second respondent:

no answer

Third respondent:

a) 1/2 to 2M dollars

b) $500 to $1,500 semiconductor type
$500 to $400 core memory type

c) + or - 25%

Fourth respondent:

a) "$1 Million/program including CPU, IOU and memory.

b) $1K for a specific card

c) 5%

Fifth respondent:

a) "In answering this question, we will assume that the reliability of the computer is good and that ECPs to modify the design to improve reliability will not be required. The fixed costs should then be approximately 80% of the total repair costs."

b) "Using the same assumptions, concerning reliability, as we did in Question 28, 20% of the total repair costs.

c) + or - 10%

Sixth respondent:

a) "Dependent on alternative specifications selected."

b) "Alternative specifications can drive costs as much as 3:1."

c) "Tied to a given requirements spec, a reasonable est + or - 10-20%."

Question 31: What are your estimates of the investment required as a function of time to field a new machine?

$0.0 to government

$3-5M over 2 years

$20M for vendor, $30M for government over 4-5 years at approximately 10%, 25%, 25%, 25%, 15% for those years

About $2M at $250K, $750K, $750K, $250K

Question 32: a) Do you believe current life cycle cost modules are adequate for the job envisioned? b) If inadequate, what are the critical aspects that must be developed?

a) 0-yes
4-no
2-no answer

b) Acquisition costs must reflect 1) savings from capturing commercial software, 2) savings from not having to pay for computer development, 3) savings from more modern software.

Updated models for "new technology designs and maintenance methods"

LCC models must be made visible to contractors on procurements

Specific contract award and contract execution incentives spelled out

Need more detail in areas of 1) man hours to make repairs, 2) percentage of repairs made at each maintenance level, 3) number of spares available, 4) MTTR, 5) MTBF.

Method must be developed to verify the credibility of the numbers used as an input to the model.

Payment of final units delayed until LCC inputs are computed with actual field data and compared to LCC originally computed with projected data; then determine additional rewards and penalties.

One further comment offered was that it would be desirable to see software compatible replacements of prior generation computers, but that DoD should participate in the feasibility and design phase.

Appendix C

ACCREDITATION CYCLE LENGTHS FOR MINIMUM LIFE-CYCLE COSTS

INTRODUCTION

This Appendix contains a cost model that shows the effect of new technology on life-cycle costs for embedded computer systems. The model uses a present value of future expenditures, thereby discounting both future costs and savings to reflect the greater value of present monies. The model also presumes that technology improvements occur on a regular basis, thereby increasing some future savings by using the most modern technology possible. On the basis of the model run over several sets of parameters, there is strong evidence that the accreditation cycle should allow for the introduction of new technology about 6 to 8 years after the initial procurement of an embedded computer system. The best possible time to introduce new technology does depend on relative sizes of R&D expenditures, procurement costs, and logistics costs. However, the 6 to 8-year accreditation cycle time appears to be optimum or having a cost not very far from optimum for a wide variety of assumptions.

The Cost Model

We assume that there are three cost components to an embedded computer system that determine its life-cycle cost. Let R denote the R&D expenditure, P the cost for procuring the system, and L be the annual logistics cost. Then:

$$LCC = R + P + 20L$$

where LCC is the life-cycle cost to develop and run the system for 20 years. The logistics component of the cost includes all annual expenditures such as spares, maintenance personnel, training of maintenance personnel, inventory,

transportation, and warehousing costs. The LCC as presented here is undiscounted, that is, all dollars are weighted equally regardless of when they are spent. To obtain a more realistic estimate of cost, we use a discount factor d that weighs the value of a dollar n years in the future as $(1 - d)^n$. For most present value calculations it is usual to use the discount factor .1 (10%), although in recent times there is strong evidence that 10% may be too low for the next two decades.

To compute a discounted logistics cost, note that expending L dollars for each of 20 years incurs a discounted cost of

$$L + (1-d)L + (1-d)^2L + \ldots + (1-d)^{19}L$$
$$= L \cdot (1 - (1-d)^{20})/d$$

Then a discounted LCC, denoted LCC_d is given by

$$LCC_d = R_d + P_d + L_d$$
$$= R_d + P_d + L(1 - (1-d)^{20})/d$$

The accreditation cycle problem is essentially the following: Given that an embedded computer system is now in the field, at what future point in time should this be replaced by new equipment in order to lower or minimize total life-cycle cost. The presumption is that the present equipment and the new equipment are functionally equal and that the new equipment will have a lower logistics cost because of advances in technology. However, by introducing the new equipment we must incur a cost for R&D and procurement which may mitigate the logistics savings. Therefore the new equipment incurs a cost of

$$\text{New costs} = R_{new} + P_{new} + L_{new}(20 - k)$$

on an undiscounted basis if the new system is introduced k years in the future. The savings by using the new system on an undiscounted basis are:

$$\text{Costs saved} = L(20 - k)$$

which are the logistics costs that would have been incurred had not the new system replaced the old. We wish to choose k so that the costs saved minus the new costs are maximized on a discounted basis. Because technology improves in time, the longer we wait to insert new technology the lower the value of L_{new}, which tends to increase the annual savings from the new equipment. However, if we wait too long before inserting new technology, then the annual savings are realized over a shorter period of time, and the life-cycle costs are not as low as they are if new equipment is introduced earlier.

To model the effect of time on technology we introduce the technology improvement factor t. Each year we assume that we can purchase equal capability in computers for a fraction t less than the cost the year before. That is, if it costs P to purchase a system today, then by using technology k years newer we can purchase that same system for $P(1-t)^k$ dollars, undiscounted, any time after k years in the future. Similarly, the logistics costs for that system due to built-in test, higher reliability, smaller size and weight, etc., will be $L(1-t)^k$ dollars per year, undiscounted.

At this point we can calculate the discounted costs and savings. First we assume that R&D is modeled at a constant cost in fixed dollars and does not decrease with technology improvement. Then R_{new} if expended k years in the

future is given in discounted dollars as:

$$R_{new}(\text{discounted}) = R_{new}(1-d)^k \qquad (1)$$

Procurement dollars benefit by advances in technology by the amount of a fraction t per year. Consequently, a procurement of a new system k years in the future has a discounted cost of

$$P_{new}(\text{discounted}) = P_{new}[(1-d)(1-t)]^k \qquad (2)$$

Logistics dollars are expended in the future at the rate of L_{new} per year, undiscounted. Because of technology advances, $L_{new} = L(1-t)^k$ if we are using a technology k years newer than the present one. Then on a discounted basis, the logistics costs for running this system for 20 - k years starting k years in the future is:

$$L_{new}(1 - (1-d)_{20-k})/d = L(1-t)^k(1 - (1-d)20-k)(1-d)^k/d \qquad (3)$$

The savings in logistics is the difference in discounted dollars of running a system at an annual cost of L dollars for 20 years, versus the cost of running the system for only k years. For the remainder of the 20-year life cycle, logistics charges are computed by the logistics formula given above. So the cost savings in logistics is

$$L(1 - (1-d)^{20})/d - L(1 - (1-d)^k)/d =$$
$$L(1-d)^k(1 - (1-d)^{20-k})/d \qquad (4)$$

The total savings in logistics from introducing new technology at year k is then the savings in equation (4) less the costs in equation (3) which yields:

$$\text{Logistics savings(discounted)} = L(1-d)^k(1 - (1-t)^k)(1 - (1-d)^{20-k})/d \qquad (5)$$

To determine suitable accreditation cycles we must determine when the logistics savings in (5) is maximized, and then balance this expenditure against discounted R&D and procurement expenditures from equations (1) and (2). Generally speaking, the discounting and the technology improvement factors force the costs in (1) and (2) lower as time increases. The optimum point for inserting new technology will be sometime after the point at which logistics savings are maximum, since the later time may achieve total lower cost from lower R&D and procurement costs in spite of the slightly reduced logistics savings.

Calculations and Discussion of Findings

The attached tables show the multiplicative factors in equation (1), (2) and (5), to cover the R&D, procurement, and logistics factors for various values of discount factor d, technology improvement factor t, and year of introduction k. The discount factors used are 5%, 10%, and 15%, and the technology improvement factors run from 5% to 25% in increments of 5%. For some aspects of technology the historic trend has been as high as 20% per year. Most notably, this occurs in memory technology. Not all aspects of device technology have shown this improvement, and there is some doubt that military technology can improve at the rate of 20% per annum in costs over long periods of time. Consequently the factors studied should bracket the

possible range of such factors over the next several years.

The primary observation is that logistics cost savings are maximized in every calculation for periods of time in the four to seven year range. For a few instances the maximum comes later, but the values in the four to seven year range are not far from optimal. The coefficient given in the logistics savings column is the multiplier of the annual cost of logistics L in the present technology. Consequently, the discounted logistics savings varies from about L/2 to 6L on a discounted basis over the 20-year life cycle, depending on the technology and discount factors.

The actual savings is somewhat less than the logistics savings because of the R_{new} and P_{new} terms that account for R&D and procurement expenditures. R&D for a system to be fielded in seven years is likely to be initiated in the present year or in the next few years. Consequently, these expenditures will be discounted by very small factors and can be found in the tables of coefficients. A procurement for a system fielded in seven years is likely to start in six to seven years, so that this expenditure is likely to have a discount and technology factor for the kth year of k-1th year if the logistics is to begin the kth year. The coefficient for discounting the procurement is found in the Purchase Costs columns of the tables.

Since savings in logistics tend to be maximized in the four to seven year time frame, total life-cycle costs will be minimized at some point in time slightly later than the optimum logistics time in order to reduce the cost of R&D and procurement. Consequently, we believe that the time frame six to eight years after initial introduction of a system appears to the time when new technology will achieve the greatest savings. These findings do depend on the relative costs of R&D, procurement, and logistics, but the findings should not alter the date of introduction by more than a few years as the ratios in

costs and savings vary over reasonable ratios.

Recommendations

It is worthwhile to continue this cost exercise to the extent that realistic values of R, and P, and L are used to determine total costs. Then we can obtain more nearly exact data on the best point to introduce new technology according to this model. Because of the nature of the data, we expect that the exact calculations will confirm the general findings of this report and substantiate that an accreditation cycle of six to eight years is reasonable to have in a 20-year program.

It is also worth some study to determine how strongly these results depend on the 20 year assumption because costs that may be incurred beyond 20 years are so heavily discounted by the model. Obviously, if the life cycle were reduced to 10 years from 20 years, the savings achievable in logistics would be greatly diminshed, especially if the new technology were introduced late in the 10-year life cycle. A shorter life cycle, say 10 years, would undoubtedly show that technology should not be introduced during a system life span, but rather the entire system should be replaced with a new one at the end of the life of the system.

For certain values of parameters the maximum savings in logistics is equal to only about half the annual logistics expenditure, with this savings spread out over the 20-year period on a discounted basis. For these values of parameters the model may indicate that new technology should not be introduced, since the savings may be about the same magnitude as the cost of the new R&D and procurement. The advisability of using new technology in such cases is largely influenced by the R&D and procurement.

DISCOUNTED COST FACTORS

Annual Discount Factor: 5%
Annual Logistics Improvement Factor: 5%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.95000	0.90250	0.59151
2	0.90250	0.81451	1.06083
3	0.85738	0.73509	1.42308
4	0.81451	0.66342	1.69178
5	0.77378	0.59874	1.87895
6	0.73509	0.54036	1.99532
7	0.69834	0.48767	2.05041
8	0.66342	0.44013	2.05269
9	0.63025	0.39721	2.00969
10	0.59874	0.35849	1.92808
11	0.56880	0.32353	1.81375
12	0.54036	0.29199	1.67193
13	0.51334	0.26352	1.50724
14	0.48767	0.23783	1.32374
15	0.46329	0.21464	1.12500
16	0.44013	0.19371	0.91417
17	0.41812	0.17482	0.69400
18	0.39721	0.15778	0.46690
19	0.37735	0.14240	0.23496
20	0.35849	0.12851	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 5%
Annual Logistics Improvement Factor: 10%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.95000	0.85500	1.18303
2	0.90250	0.73103	2.06725
3	0.85738	0.62503	2.70398
4	0.81451	0.53440	3.13651
5	0.77378	0.45691	3.40135
6	0.73509	0.39066	3.52924
7	0.69834	0.33401	3.54603
8	0.66342	0.28558	3.47340
9	0.63025	0.24417	3.32953
10	0.59874	0.20877	3.12961
11	0.56880	0.17850	2.88631
12	0.54036	0.15261	2.61015
13	0.51334	0.13048	2.30988
14	0.48767	0.11156	1.99270
15	0.46329	0.09539	1.66454
16	0.44013	0.08156	1.33025
17	0.41812	0.06973	0.99378
18	0.39721	0.05962	0.65831
19	0.37735	0.05097	0.32638
20	0.35849	0.04358	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 5%
Annual Logistics Improvement Factor: 15%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.95000	0.80750	1.77454
2	0.90250	0.65206	3.01928
3	0.85738	0.52654	3.85018
4	0.81451	0.42518	4.35950
5	0.77378	0.34333	4.62053
6	0.73509	0.27724	4.69138
7	0.69834	0.22387	4.61806
8	0.66342	0.18078	4.43685
9	0.63025	0.14598	4.17637
10	0.59874	0.11788	3.85903
11	0.56880	0.09518	3.50239
12	0.54036	0.07686	3.12008
13	0.51334	0.06207	2.72266
14	0.48767	0.05012	2.31825
15	0.46329	0.04047	1.91300
16	0.44013	0.03268	1.51158
17	0.41812	0.02639	1.11741
18	0.39721	0.02131	0.73302
19	0.37735	0.01721	0.36015
20	0.35849	0.01389	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 5%
Annual Logistics Improvement Factor: 20%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.95000	0.76000	2.36606
2	0.90250	0.57760	3.91690
3	0.85738	0.43898	4.86916
4	0.81451	0.33362	5.38469
5	0.77378	0.25355	5.58422
6	0.73509	0.19270	5.55762
7	0.69834	0.14645	5.37159
8	0.66342	0.11130	5.07550
9	0.63025	0.08459	4.70576
10	0.59874	0.06429	4.28909
11	0.56880	0.04886	3.84497
12	0.54036	0.03713	3.38752
13	0.51334	0.02822	2.92686
14	0.48767	0.02145	2.47015
15	0.46329	0.01630	2.02236
16	0.44013	0.01239	1.58686
17	0.41812	0.00942	1.16583
18	0.39721	0.00716	0.76061
19	0.37735	0.00544	0.37192
20	0.35849	0.00413	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 5%
Annual Logistics Improvement Factor: 25%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.95000	0.71250	2.95757
2	0.90250	0.50766	4.76012
3	0.85738	0.36171	5.76841
4	0.81451	0.25771	6.23465
5	0.77378	0.18362	6.33487
6	0.73509	0.13083	6.19156
7	0.69834	0.09322	5.88973
8	0.66342	0.06642	5.48813
9	0.63025	0.04732	5.02716
10	0.59874	0.03372	4.53443
11	0.56880	0.02402	4.02863
12	0.54036	0.01712	3.52226
13	0.51334	0.01220	3.02354
14	0.48767	0.00869	2.53774
15	0.46329	0.00610	2.06810
16	0.44013	0.00441	1.61645
17	0.41812	0.00314	1.18372
18	0.39721	0.00224	0.77020
19	0.37735	0.00160	0.37576
20	0.35849	0.00114	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 10%
Annual Logistics Improvement Factor: 5%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.90000	0.85500	0.38921
2	0.81000	0.73103	0.67121
3	0.72900	0.62503	0.86634
4	0.65610	0.53440	0.99151
5	0.59049	0.45691	1.06077
6	0.53144	0.39066	1.08576
7	0.47830	0.33401	1.07609
8	0.43047	0.28558	1.03966
9	0.38742	0.24417	0.98296
10	0.34868	0.20877	0.91128
11	0.31381	0.17850	0.82891
12	0.28243	0.15261	0.73934
13	0.25419	0.13048	0.64536
14	0.22877	0.11156	0.54917
15	0.20589	0.09539	0.45252
16	0.18530	0.08156	0.35678
17	0.16677	0.06973	0.26298
18	0.15009	0.05962	0.17190
19	0.13509	0.05097	0.08411
20	0.12158	0.04358	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 10%
Annual Logistics Improvement Factor: 10%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.90000	0.81000	0.77842
2	0.81000	0.65610	1.30800
3	0.72900	0.53144	1.64612
4	0.65610	0.43047	1.83823
5	0.59049	0.34868	1.92025
6	0.53144	0.28243	1.92046
7	0.47830	0.22877	1.86102
8	0.43047	0.18530	1.75923
9	0.38742	0.15009	1.62850
10	0.34868	0.12158	1.47916
11	0.31381	0.09848	1.31909
12	0.28243	0.07977	1.15423
13	0.25419	0.06461	0.98902
14	0.22877	0.05233	0.82669
15	0.20589	0.04239	0.66955
16	0.18530	0.03434	0.51917
17	0.16677	0.02781	0.37658
18	0.15009	0.02253	0.24238
19	0.13509	0.01825	0.11684
20	0.12158	0.01478	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 10%
Annual Logistics Improvement Factor: 15%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.90000	0.76500	1.16764
2	0.81000	0.58523	1.91037
3	0.72900	0.44770	2.34389
4	0.65610	0.34249	2.55499
5	0.59049	0.26200	2.60854
6	0.53144	0.20043	2.55284
7	0.47830	0.15333	2.42364
8	0.43047	0.11730	2.24721
9	0.38742	0.08973	2.04270
10	0.34868	0.06865	1.82391
11	0.31381	0.05251	1.60065
12	0.28243	0.04017	1.37973
13	0.25419	0.03073	1.16577
14	0.22877	0.02351	0.96175
15	0.20589	0.01790	0.76949
16	0.18530	0.01376	0.58994
17	0.16677	0.01053	0.42343
18	0.15009	0.00805	0.26988
19	0.13509	0.00616	0.12893
20	0.12158	0.00471	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 10%
Annual Logistics Improvement Factor: 20%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.90000	0.72000	1.55685
2	0.81000	0.51840	2.47832
3	0.72900	0.37325	2.96423
4	0.65610	0.26874	3.15583
5	0.59049	0.19349	3.15260
6	0.53144	0.13931	3.02421
7	0.47830	0.10031	2.81911
8	0.43047	0.07222	2.57067
9	0.38742	0.05200	2.30163
10	0.34868	0.03744	2.02717
11	0.31381	0.02696	1.75721
12	0.28243	0.01941	1.49799
13	0.25419	0.01397	1.25320
14	0.22877	0.01006	1.02477
15	0.20589	0.00724	0.81348
16	0.18530	0.00522	0.61932
17	0.16677	0.00376	0.44177
18	0.15009	0.00270	0.28004
19	0.13509	0.00195	0.13314
20	0.12158	0.00140	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 10%
Annual Logistics Improvement Factor: 25%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.90000	0.67500	1.94606
2	0.81000	0.45563	3.01185
3	0.72900	0.30755	3.51167
4	0.65610	0.20759	3.65397
5	0.59049	0.14013	3.57638
6	0.53144	0.09459	3.36917
7	0.47830	0.06384	3.09104
8	0.43047	0.04310	2.77967
9	0.38742	0.02909	2.45883
10	0.34868	0.01964	2.14313
11	0.31381	0.01325	1.84115
12	0.28243	0.00895	1.55758
13	0.25419	0.00604	1.29459
14	0.22877	0.00408	1.05281
15	0.20589	0.00275	0.83188
16	0.18530	0.00186	0.63087
17	0.16677	0.00125	0.44855
18	0.15009	0.00085	0.28357
19	0.13509	0.00057	0.13451
20	0.12158	0.00039	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 15%
Annual Logistics Improvement Factor: 5%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.85000	0.80750	0.27041
2	0.72250	0.65206	0.44443
3	0.61413	0.52654	0.54708
4	0.52201	0.42518	0.59759
5	0.44371	0.34333	0.61071
6	0.37715	0.27724	0.59762
7	0.32058	0.22387	0.56676
8	0.27249	0.18078	0.52446
9	0.23162	0.14598	0.47539
10	0.19687	0.11788	0.42297
11	0.16734	0.09518	0.36964
12	0.14224	0.07686	0.31710
13	0.12091	0.06207	0.26651
14	0.10277	0.05012	0.21863
15	0.08735	0.04047	0.17387
16	0.07425	0.03268	0.13247
17	0.06311	0.02639	0.09447
18	0.05365	0.02131	0.05982
19	0.04560	0.01721	0.02839
20	0.03876	0.01389	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 15%
Annual Logistics Improvement Factor: 10%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.85000	0.76500	0.54083
2	0.72250	0.58523	0.86607
3	0.61413	0.44770	1.03949
4	0.52201	0.34249	1.10792
5	0.44371	0.26200	1.10553
6	0.37715	0.20043	1.05704
7	0.32058	0.15333	0.98017
8	0.27249	0.11730	0.88745
9	0.23162	0.08973	0.78760
10	0.19687	0.06865	0.68656
11	0.16734	0.05251	0.58822
12	0.14224	0.04017	0.49504
13	0.12091	0.03073	0.40844
14	0.10277	0.02351	0.32911
15	0.08735	0.01799	0.25726
16	0.07425	0.01376	0.19277
17	0.06311	0.01053	0.13528
18	0.05365	0.00805	0.08435
19	0.04560	0.00616	0.03944
20	0.03876	0.00471	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 15%
Annual Logistics Improvement Factor: 15%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.85000	0.72250	0.81124
2	0.72250	0.52201	1.26492
3	0.61413	0.37715	1.48013
4	0.52201	0.27249	1.53993
5	0.44371	0.19687	1.50179
6	0.37715	0.14224	1.40511
7	0.32058	0.10277	1.27649
8	0.27249	0.07425	1.13361
9	0.23162	0.05365	0.98792
10	0.19687	0.03876	0.84657
11	0.16734	0.02800	0.71377
12	0.14224	0.02023	0.59175
13	0.12091	0.01462	0.48143
14	0.10277	0.01056	0.38288
15	0.08735	0.00763	0.29566
16	0.07425	0.00551	0.21904
17	0.06311	0.00398	0.15211
18	0.05365	0.00288	0.09392
19	0.04560	0.00208	0.04352
20	0.03876	0.00150	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 15%
Annual Logistics Improvement Factor: 20%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.85000	0.68000	1.08165
2	0.72250	0.46240	1.64098
3	0.61413	0.31443	1.87186
4	0.52201	0.21381	1.90206
5	0.44371	0.14539	1.81502
6	0.37715	0.09887	1.66455
7	0.32058	0.06723	1.48477
8	0.27249	0.04572	1.29678
9	0.23162	0.03109	1.11315
10	0.19687	0.02114	0.94092
11	0.16734	0.01437	0.78359
12	0.14224	0.00977	0.64247
13	0.12091	0.00665	0.51753
14	0.10277	0.00452	0.40797
15	0.08735	0.00307	0.31257
16	0.07425	0.00209	0.22995
17	0.06311	0.00142	0.15870
18	0.05365	0.00097	0.09746
19	0.04560	0.00066	0.04494
20	0.03876	0.00045	0.00000

DISCOUNTED COST FACTORS

Annual Discount Factor: 15%
Annual Logistics Improvement Factor: 25%

Yr	R&D Costs	Purchase Costs	Logistics Savings
1	0.85000	0.63750	1.35207
2	0.72250	0.40641	1.99424
3	0.61413	0.25908	2.21755
4	0.52201	0.16517	2.20230
5	0.44371	0.10529	2.05900
6	0.37715	0.06712	1.85443
7	0.32058	0.04279	1.62800
8	0.27249	0.02728	1.40221
9	0.23162	0.01739	1.18918
10	0.19687	0.01109	0.99474
11	0.16734	0.00707	0.82102
12	0.14224	0.00451	0.66803
13	0.12091	0.00287	0.53463
14	0.10277	0.00183	0.41913
15	0.08735	0.00117	0.31964
16	0.07425	0.00074	0.23424
17	0.06311	0.00047	0.16114
18	0.05365	0.00030	0.09869
19	0.04560	0.00019	0.04541
20	0.03876	0.00012	0.00000

Appendix D

TECHNOLOGY INSERTION IN MILITARY COMPUTER SYSTEMS

TECHNOLOGY INSERTION IN MILITARY COMPUTER SYSTEMS

The two questions to be studied are essentially:

1. How can we determine if it is cost-effective to introduce new technology into embedded computers at some specific time?
2. With what frequency should we plan to insert new technology into embedded computers?

The first question is posed in such a way as to require a simple yes/no answer, that is, should a new system be introduced now, or should the old system be continued for a longer period of time? The second question requires a more extensive computational model because it admits to solutions that indicate at what specific times new technology should be introduced.

In any event, the key to introducing new technology is that the future benefits of the new technology are far greater than the present costs of introducing it. To the extent that benefits outweigh the costs, it is worthwhile to introduce new technology. However, if gains are small and introduction costs are high, it is better to retain older technology. The rate of introduction of new technology cannot be too fast because new systems must be installed and stable for a number of years in order to derive some cost benefit.

The first step in evaluating cost benefit and return on investment is to incorporate the time value of money into the model. That is, future costs and future savings have to be discounted by a constant annual factor (10% is the typical figure for this type of study). Discounting tends to weigh near-term gains more heavily than long-term gains and to penalize near-term costs. The formula for the time value of future money reflected back to the present is:

$$V(\text{year}) = V(\text{year} + N)\ (1-D)^N$$

where V(year) is the value at time "year," V(year + N) is the value N years in the future, and D is the discount factor measured as a fraction. (D is .1 for a 10% discount.)

The easier question to answer is the question concerning whether or not to deploy a given new technology today. A cost analysis should yield the following information:

1. Investment dollars (R & D) required to field the new technology.

2. Reductions in cost from using the new technology. Reductions may be in terms of volume, weight, power, maintenance, spare parts, purchase price, software system costs, or other similar factors.

As an example of a calculation of costs and benefits, consider the life-cycle cost model described by Stone in Ref. 1. This model breaks down computer costs into three areas:

1. Common costs. These are the fixed costs for hardware and software R & D, product specification, product planning, and the development of basic support software such as compilers and operating systems. Each common cost is incurred once, regardless of the number of systems procured. These costs would not be incurred if old systems were left in place and the introduction of new technology were deferred.

2. Hardware costs. These are the variable costs that are proportional to the number of computers procured. These costs include hardware logistics support costs.

3. Software costs. These costs are incurred once per application. Since software replication costs are negligible, working software can be copied for all sites for essentially nothing after the investment in producing the working software has been made.

In this model the future benefits lie in the potentially large reduction in hardware and software costs in the future if the present investment in dollars is made. Future hardware is assumed to benefit from new semiconductors which are less expensive per function and inherently more reliable than the parts they replace. Software savings may be achieved in the future by making use of advanced software tools that may exist for the new computer system and are not available on the one already deployed. With more powerful software tools, the hope is that new software systems can be implemented and maintained at much lower cost than if those new systems were implemented with the tools available at present.

It is not difficult to do a present value calculation on the potential future costs and compare this value to the present value of the investment dollars required to field the new technology. To do so, simply apply the formula for the time value of money to reflect all costs and all savings in terms of 1980 dollars.

A very large savings should lend weight to a decision to make the change in technology. A moderate to low savings indicates increased risk in achieving the desired savings. It may be necessary to defer the introduction of new technology for a year or two, at which time the savings from the most recent advances may be much greater and lend more support to the introduction of new technology.

The more difficult question concerns how frequently should new technology be introduced. Here the model presumes that each year there are some new gains available because of the advances of the past year. There are various methods by which one can estimate when to introduce new technology. We propose one such method here.

The basic idea is to estimate annual expenditures over one or two full life cycles of equipment. The expenditures depend on exactly what technology is used during each year of the period studied. For example, with present technology used over the entire time span, let us assume a fixed annual expenditure of 1 unit. If we assume a discount factor of 10%, then we expend 1 unit the first year, .90 units the second, .81 the third, etc., and the present value of these expenditures over a 20-year period is 7.91. There is a simple formula we can use to compute the sum of future values reflected back to the present, because such a sum is a geometric sequence. Let F be the annual expenditure, N be the number of years over which the expenditure occurs, and PV (N) be the value of those expenditures reflected back to the present. Then:

$$PV(N) = F + (1-D)F + (1-D\_\ F + \ldots + (1-D)^{N-1}F$$

$$PV(N) = F\ [1 - (1-D)^{N}/D]$$

For an annualized expenditure normalized to a unit cost, the factor F is 1. Now let us compare this cost to a scenario that calls for the introduction of new technology sometime in the future. Suppose we incur development costs of C in five years, and in 10 years we reduce annual expenditures to 1/2 when the new technology replaces the old. Then the cost equation for this example becomes:

cost = old technology for 10 years + development costs + new technology for N - 10 years.

$$= [1 - (1-D)^{10}/D] + C(1-D)^{5} + (1-D)^{10}[1 - (1-D)^{N-10}/2D]$$

In this equation the first term represents the cost of old technology run for 10 years, as evidenced by the exponent of 10. The next term has the coefficient C,

which is the development cost of the new technology, and it is reflected back to its present value from 5 years in the future. The final term is the cost of the new technology fielded for N-10 years at half the annual cost (note the factor of 2 in the denominator of this term). This savings is reflected back to the present value by multiplying by (1-D) raised to a power.

To model when and how often to insert new technology we can use the formulation given here with the following additional information:

1. We must have an estimate for the investment required to field the new technology. This may depend on time in that R & D expenditures tend to become more expensive as one attempts to wring greater benefits from technology.

2. What are the projected annual savings from new technology? Again this is time dependent in that the more advanced technology brings with it greater benefits.

3. Over how long a period of time are the expenditures made?

A simple computational approach will suffice to answer the questions after the basic input data have been established. The best approach is to look for the optimum point to change technology exactly once over the time span. This gives some information as to where the benefits of new technology are greatest. Then a more complex optimization can be performed to determine where to change technology exactly twice over the length of time. This requires some computational search that is basically a dynamic programming exercise and is quite easy to perform. If necessary, the optimization can be carried to the point of determining where to change technology exactly three times, for which the computation becomes more difficult. If the time period is on the order of 20 years, we expect that the model will show that the benefits of new technology are lost if changes become more frequent than a few times.

REFERENCES

3. Stone, H. "Life-cycle cost analysis of instruction-set architecture standardization for military computer systems," Computer, Vol. 12, No. 4, pp. 35-47, April 1979.

Appendix E
COSTS OF IMPLEMENTING LOGISTICS SYSTEMS

COSTS OF IMPLEMENTING LOGISTICS SYSTEMS

NUMBER OF VENDORS VS. LOGISTICS COSTS

In 1978 and 1979 a working group consisting of representatives of all three services attempted to put together a life-cycle cost model of logistics support that would show the effects of standardization of logistics costs. (Ref. 1) The issues studied by this cost model were principally the effect of standardization on standard cards, standard modules, or standard chassis for computer systems, with each standard specified to form, fit, and function. This study also attempted to incorporate the effects of multiple suppliers on costs. Navy participation in the study was very limited because of a lack of funding, so that the ultimate report was issued with inputs only from Air Force and Army representatives.

The findings in the report are qualitative, not quantitative, because of the inability to obtain sufficient data within the time and funding constraints of the study. The principal factors in logistics scenarios as determined by that report are repeated here, and we examine these factors to determine how each is affected by multiple suppliers of standard parts. We also estimate the major cost contributions from having many versus few suppliers to isolate the factors that strongly influence costs.

A second study funded solely by the Army in 1979 attempted to quantify the cost of spare computer components as a function of logistic support concepts. (Ref. 2) That study also modeled the effects of multiple suppliers on spares costs. In this report we repeat the study for Navy data to show how spares costs vary with the number of suppliers.

FINDINGS OF THE 1978-1979 STUDY

The pertinent factors that contribute to logistics costs were identified as the following:

1. Contractor support costs. These are costs of warranty, repair, and resupply paid to contractors for maintenance of equipment. The costs are typically paid to the original supplier of the equipment. In some instances, the maintenance contractor is distinct from the original contractor.

2. Inventory (pipeline and float). These costs pay for equipment that is not directly in use because it is either in transit (pipeline) or in storage (float) where it is not directly available for use.

3. Transportation. These costs are incurred for shipping failed units back to the point of repair and for shipping replacement units back to the supply center.

4. Repair parts. These costs cover the purchase of extra parts that have to be held in order to repair or replace failed equipment. Included in these costs are the costs of spares held at or near the site at which equipment is used, plus spares held at repair and resupply depots.

5. Personnel, training, and facilities. These costs are for service personnel to be trained in the repair and replacement of failed equipment. The personnel may be military or civilian, but in either case are direct employees of the government so that the government incurs the costs of training and equipping the personnel.

6. Specifications, documentation, technical manuals, test and diagnostic equipment. These costs are incurred for each item procured by the military. However, the magnitude of the costs depends on whether the military is directly responsible for the repair of the equipment or if the equipment is to be sent back to a contractor for repair under warranty or contract. If the military takes over repairs, these costs are much greater than if the repairs are done under warranty.

There are two possible prevailing methods for implementing logistics system:

1. Contractor repair. Isolate failures to a replaceable module, and then send failed modules back to the contractor for repair. (If the module is sufficiently inexpensive, it may be treated as an expendable item and discarded rather than repaired.)

2. In-service repair. Failed units are sent to a military repair depot where the military, rather than a contractor, performs the repair. The repair depot typically disassembles the field replaceable unit into subassemblies, identifies the failed subassembly, and performs the repair by replacing subassemblies. The subassemblies may be returned to a contractor for repair, discarded, or may be further disassembled for testing and repair depending on the cost and complexity of the units, assembles, and subassemblies.

Both of these logistics methods have specific advantages for particular situations, so that both methods are in present use. Military repair is projected to become more difficult over time due to the high costs for training personnel and the difficulty in finding personnel with the right kinds of skills who will remain in the military over a long period of time.

Given the two methodologies and the principal cost factors, let us make qualitative estimates on the effects that the number of suppliers will have on these logistic costs. To do so, we treat each of the logistics methodologies separately.

RELATIVE COST FACTORS FOR CONTRACTOR-SUPPORTED LOGISTICS

The major cost factors that contribute to contractor support as identified in Ref. 1 are:

1. Contractor support
2. Inventory (pipeline and float)
3. Specifications and documentation
4. Transportation

The contractor support costs are dominant because repairs are not attempted within the military. Hence, this reduces the size of personnel costs, training, facilities, and all other costs related to repair in the service.

The question we must investigate is how the principal costs vary with the number of suppliers. Let us deal with each of these factors in turn.

A. Contractor Support Costs

To estimate the effect of multiple suppliers on contractor support charges, let us first assume that the government pays for all real costs incurred by the contractor, and then estimate the costs incurred per contractor. Costs for logistics break down into two terms:

$$\text{contractor costs} = \text{fixed cost} + (\text{per item charge})(\text{number of items})$$

The fixed charges include the contractor costs for test equipment, training, documentation, etc., that are necessary to perform logistics repair. Costs for documentation and specifications delivered to the military are treated as a separate cost item.

If we assume that there is a single contractor who incurs a fixed cost F and a variable cost V to repair an item, then to repair N items the government pays:

Total contractor costs = F + VxN

If there are K contractors, each incurring a fixed cost F and a variable cost of V per item, and if each contractor repairs N/K items, the cost per contractor is:

Cost per contractor = F + Vx(N/K)

so that the total cost for all contractors is

Total contractor costs (k contractors) = KxF + V

The added cost of the extra contractors is felt in this case through the extra fixed costs incurred per contractor. The government is in a sense paying for N sets of specifications, test equipment stations, facilities, training, etc., instead of a single set. However, the formula is just a first approximation of the cost change because:

1. The use of multiple contractors will tend to introduce competition and lead to some efficiencies because the contractors will have an incentive to seek ways of reducing costs.

2. The lower volume of work per contractor may in turn introduce inefficiencies. Personnel observe fewer instances of each failure type and may require longer average times for repair. Test equipment requires a greater percentage of time for calibration if it is used less frequently. Similarly, other costs tend to increase.

The variable costs reflect the cost for personnel, repair parts, and other similar costs that grow in direct proportion to the number of items repaired. The fixed charges multiply with the number of contractors since each contractor has to have all that is necessary to run a repair facility including the personnel, training, test equipment, etc.

The principal unknowns to be determined are:

1. Fixed contractor charges.

2. As the number of suppliers increase, estimate the reduction in variable charges because of competitive pressures.

3. As the number of suppliers increase, estimate the effects of lower production volumes per supplier on the variable charges.

To obtain the necessary data for a model, it is sufficient to study two or three military embedded computer systems and obtain some idea of the magnitude of the fixed charges. There is a good deal of data for systems supplied from multiple vendors available through the Navy SEM (Standard Electronic Modules) program administered at Crane, Indiana. The data may be sufficient to model how variable costs increase or decrease with the number of suppliers since typical SEM modules start with two suppliers on initial procurement and additional suppliers become qualified in later years. The SEM modules tend to be smaller, less functional, and less expensive than modules that go into embedded computers like the UYK-7 and UYK-20. Just to be sure that the SEM data can be scaled to be indicative of the costs associated with embedded computers, we should also gather fixed and variable cost data for a computer like the UYK-20.

B. Inventory (Pipeline and Float)

The pipeline charges do not vary with the number of contractors, as they include the costs for the items in transit, which presumably will not depend on the number of contractors. Inventory charges include some storage at each contractor, which is in effect an additional supply depot for spare parts. The costs of the material on inventory here does vary with the number of suppliers, but should be small compared with the number of spares at resupply depots close to the deployment sites of the various items. Consequently, these costs can be picked up as part of a spares parts charge.

C. Specifications and Documentation

These costs cover the cost of specifications that all contractors must meet and do not include the additional costs per contractor for the contractor's own specifications and documentation.

To model how these costs vary with the number of contractors, consider a single-vendor model as a baseline, and observe how the costs change as the number of vendors increase. Typically, the specifications become stricter to assure greater interchangeability, while the documentation becomes less detailed since individual vendors will develop internal documentation for their specific designs. Therefore, the key cost factor is the cost of detailed specifications that accurately reflect how to qualify each item.

A suitable source for information is the Navy SEM program, and data developed by the SEM people over time indicate that about $30,000 to $50,000 is sufficient to develop specifications for a SEM module. Because SEM modules are much smaller than modules in minicomputers and midicomputers, specification costs may run to $100,000 for a larger module, and to several times this for a chassis. The specifications charges are incurred when two or more suppliers build to the same specifications, and do not necessarily increase as the number of suppliers increases above two. While specifications charges are lower for systems provided by a single supplier, there are still some costs for specifications incurred. To model specifications charges, it will be necessary to obtain data for systems in use today by the Navy.

D. Transportation

These charges do not vary with the number of suppliers and can be treated as a large fixed cost in a model that treats cost variations due to the number of suppliers.

RELATIVE COST FACTORS FOR IN-SERVICE (MILITARY) REPAIR

Ref. 1 lists the major cost factors for this type of logistics system. They are:

1. Personnel, training, and facilities
2. Specifications, documentation, technical manuals, test and diagnostic equipment
3. Inventory (pipeline and float)
4. Repair parts
5. Transportation

We examine each of these factors below to determine how the factors vary with the number of suppliers, and indicate what additional data should be developed to quantify this relation.

A. Personnel, Training and Facilities

These appear to be the dominant costs for repair that is done within the military. It is not certain how they vary as the number of suppliers increases because of the effect of built-in test facilities within the embedded computers.

As a general rule, built-in test reduces the skill requirements of the maintenance technician. An ideal built-in test capability permits a technician to repair modules from any supplier with equal ease and without training that is specialized to a particular supplier. In the absence of built-in test, technicians must receive separate training for equipment from each separate supplier, and costs tend to grow linearly with the number of suppliers. Actual practice will find costs somewhere between the two extremes. It will be necessary to investigate current practice to determine what these costs are and to find out the effects of built-in test systems on these costs. One way to develop the data is to interview people connected with the UYK-7 and UYK-20 computer programs and learn how they train service men for maintenance of these computers. This may be compared with the later generation AYK-14 to find if the advanced technology has had any effect on reducing personnel costs.

B. Specifications, Documentation, Technical Manuals, Test and Diagnostic Equipment

Some of these costs appear in an earlier section as contractor costs. The cost burden shifts to the military when the military undertakes the repair process. Consequently, data developed for the contractor repair model will support the military repair model. The cost burden should be about equal for the two models, except for some savings achieved for military repair by reducing proliferation of test equipment among several contractors and sharing of test equipment over several different systems. There are inefficiencies in the military repair model as well that may more than compensate for the efficiencies. Military people have less technical training and experience than people normally doing contractor repair. Consequently, the cost of the documentation and test equipment may be somewhat higher for the military than for the contractor.

We suspect that the costs for documentation and test equipment grows linearly or almost linearly with the number of suppliers, and recommend that hard data on these costs be obtained from interviews with the military.

C. Inventory (Pipeline and Float)

The costs for inventory of parts is substantially the same for the contractor repair and military repair models, with the costs being slightly less for military repair.

D. Repair Parts

A very detailed analysis of costs for spares leads to the conclusion that the cost of spares used for repairs increases with the number of suppliers, but is at most a small fraction of the total cost of spares. The model shows that spares costs for five suppliers may be 10 to 50 percent higher than spares costs for one supplier depending on deployment strategies, failure rates, and other factors.

E. Transportation

Transportation charges do not vary with the number of suppliers, but represent a large, fixed cost burden that must be included in the cost model for a logistics supply system. Transportation charges for military repair will be somewhat smaller than transportation for contractor repair, because failed units are not necessarily returned all the way to the contractor for repair if they can be repaired at depots much closer to their point of deployment.

SUMMARY

The discussion above isolates the principal cost factors for contractor and in-service repair strategies. Of these cost factors, the cost for specifications, documentation, test equipment, etc., is the cost area in which the multiple suppliers costs are felt the strongest. These costs are direct costs for both the contractor and in-service repair strategies, and they are also reflected indirectly in the contractor charges for contractor repair where they cover in-house costs for items that are not deliverables. A second area in which costs depend on the number of suppliers is the cost of spares. These costs grow slowly with the number of suppliers which indicates that the multiple supplier cost burden is more likely to be felt in terms of documentation and test equipment than in other factors. The impact of multiple suppliers on personnel costs is strongly dependent on the effectiveness of built-in test equipment. If built-in test equipment is very effective, then personnel costs may be largely independent of the number of suppliers. Otherwise, personnel costs can become proportional to the number of suppliers when repair is done in the military, and this could be a significant cost burden.

REFERENCES

1. Kohler, W., H. S. Stone, L. Bragg, and H. Hylton, "Life-Cycle Cost Factors for MCF Logistics Scenarios," report prepared for CORADCOM, Ft. Monmouth, N.J., December 1978.

2. Stone, H. and W. Kohler, "An Analysis of the Effect of Computer Standardization on Spares Inventory," report prepared for CORADCOM, Ft. Monmouth, N.J., February 1979.